Practical
Electronics

The UK's premier electronics, computing and maker magazine

Max's Cool Beans
Playing games by tilting our console

Teach-In 12
Part 9: RF sources and oscillators

Audio Out
How symmetricisers boost radio volume

Circuit Surgery
Sending & receiving balanced audio

RP2350B Dev Board: A more powerful Pico 2

Audio Out: a flexible all-pass filter system

The SmartProbe: make measurements more easily

VOLTAGE
+0.000 V
BOOT
RESET

PLUS!
Techno Talk – the history of PCB design and track routing
The Fox Report – balcony solar may be coming to the UK
Mains Earthing – why it's needed & how it varies worldwide

Jul 2026 £5.99
07
9 772632 573030

www.electronpublishing.com | @practicalelec | practicalelectronics

PCBWay

Beyond Standard —
Discover PCBWay's
Advanced PCB Lineup

About PCBWay:

With dedicated workshops for advanced PCBs, PCBWay offers full-spec, high-precision, and large-scale production. We ensure customers are supplied with PCBs that meet stringent quality standards at the lowest cost.

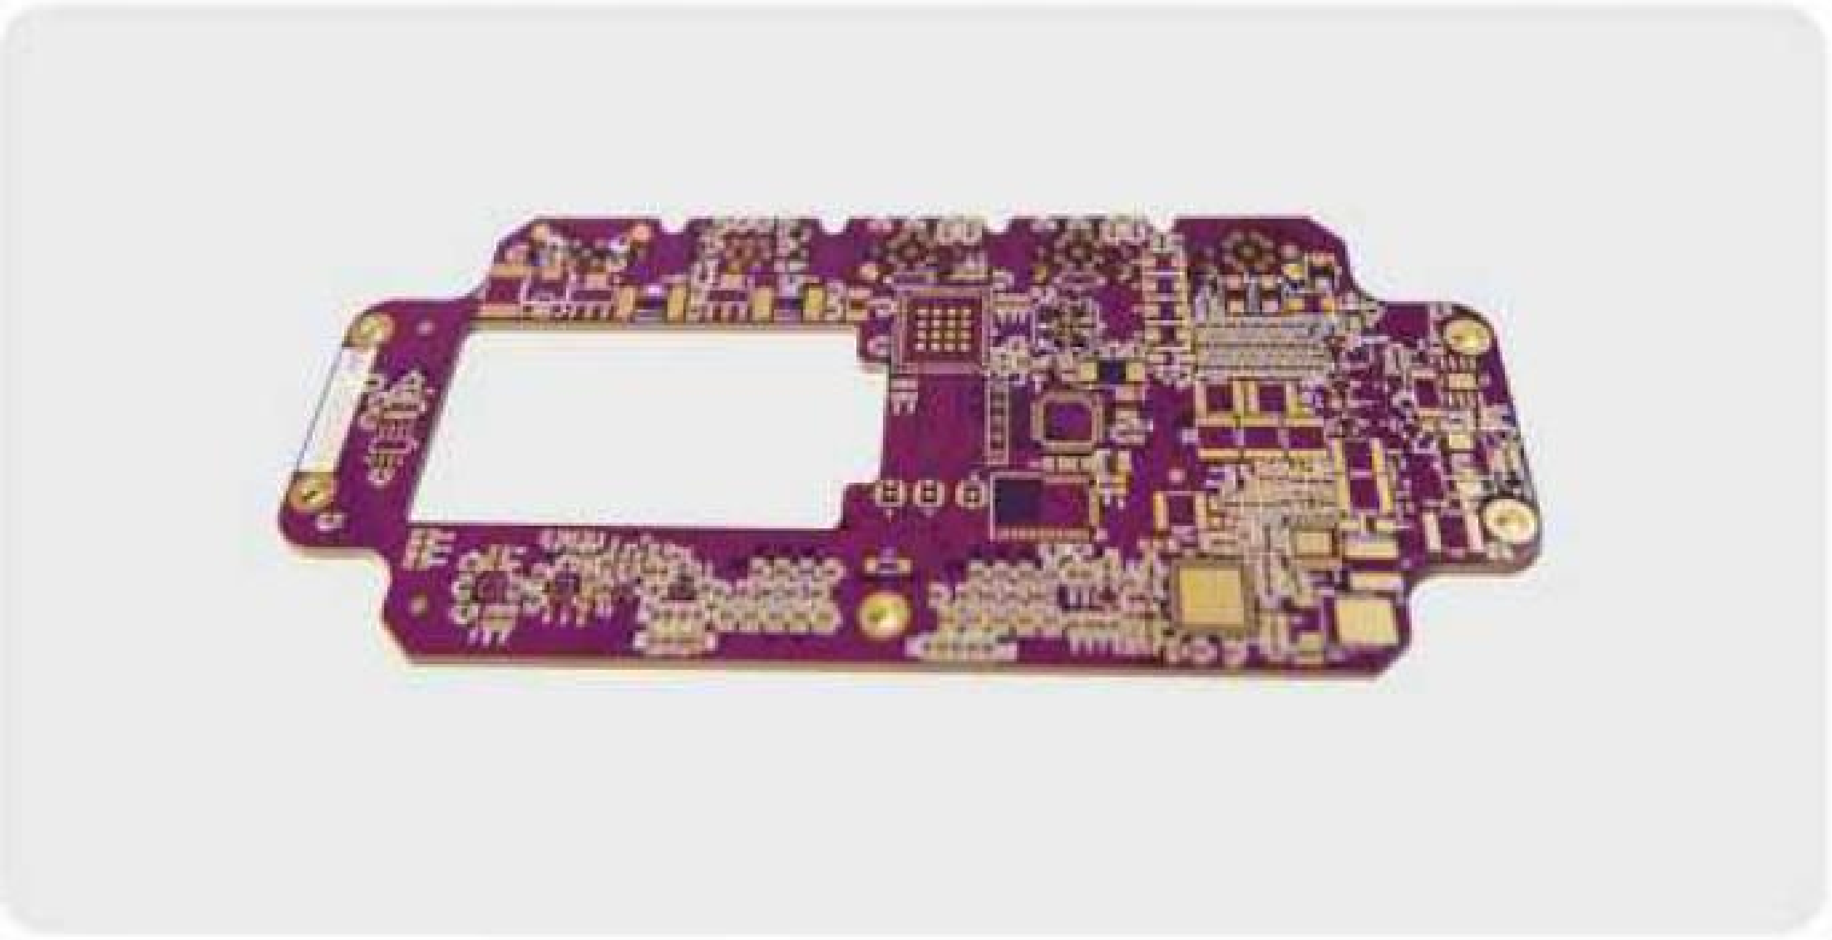

HDI PCB //

- 4-60 layers
- FR-4, High Tg, Halogen-free, PTFE
- Finished Copper: 1/3OZ–8OZ
- Min Track/Spacing (inner): 2/2mil (H/HOZ base copper)
- Min BGA Pitch: 0.35mm

High Frequency PCB //

- RO4003C, RO4350B, Ro3003, Ro3010, RT5880
- Min. Track/Spacing: 2mil
- Min. Hole Size: 0.15mm
- Finished Cooper: 1-2oz
- Board Thickness: 0.2–3.2mm

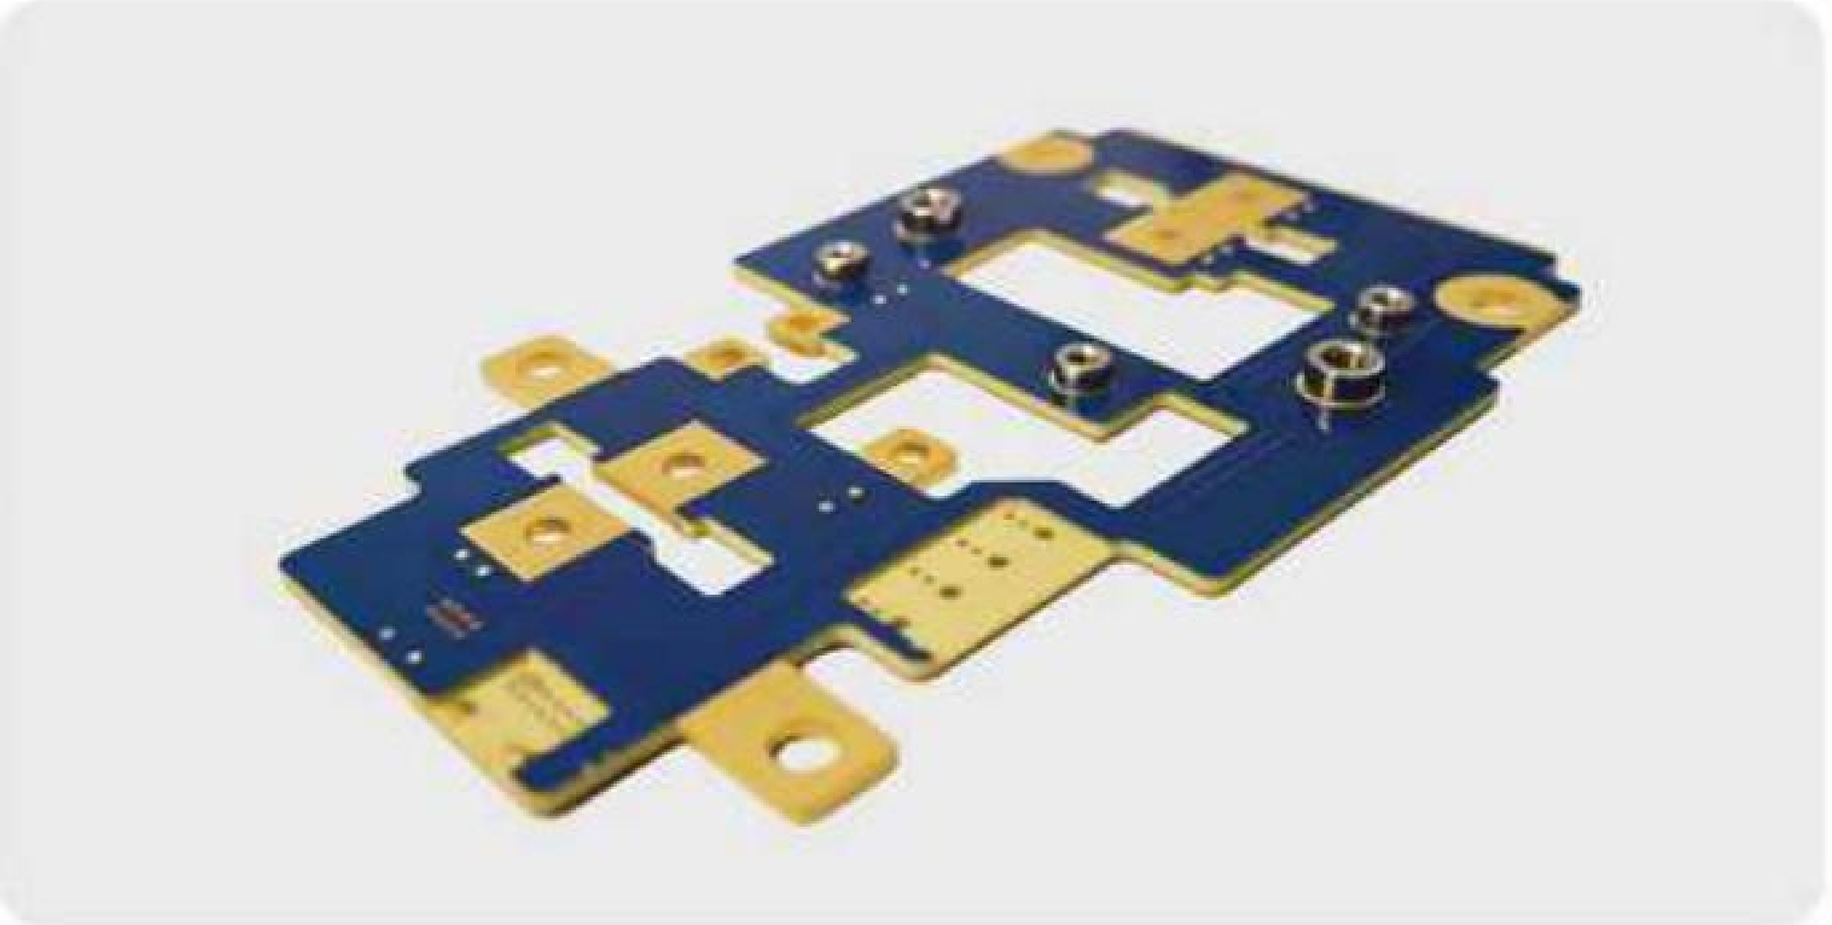

Other PCBs with Special Process //

- Embedded Copper PCB
- Backplane PCB
- Cavity PCB
- White/Partially Hybrid PCB
- Black Core PCB (After-dark)

www.pcbway.com service@pcbway.com

Practical Electronics

Volume 55. No. 07
July 2026
ISSN 2632 573X

Contents

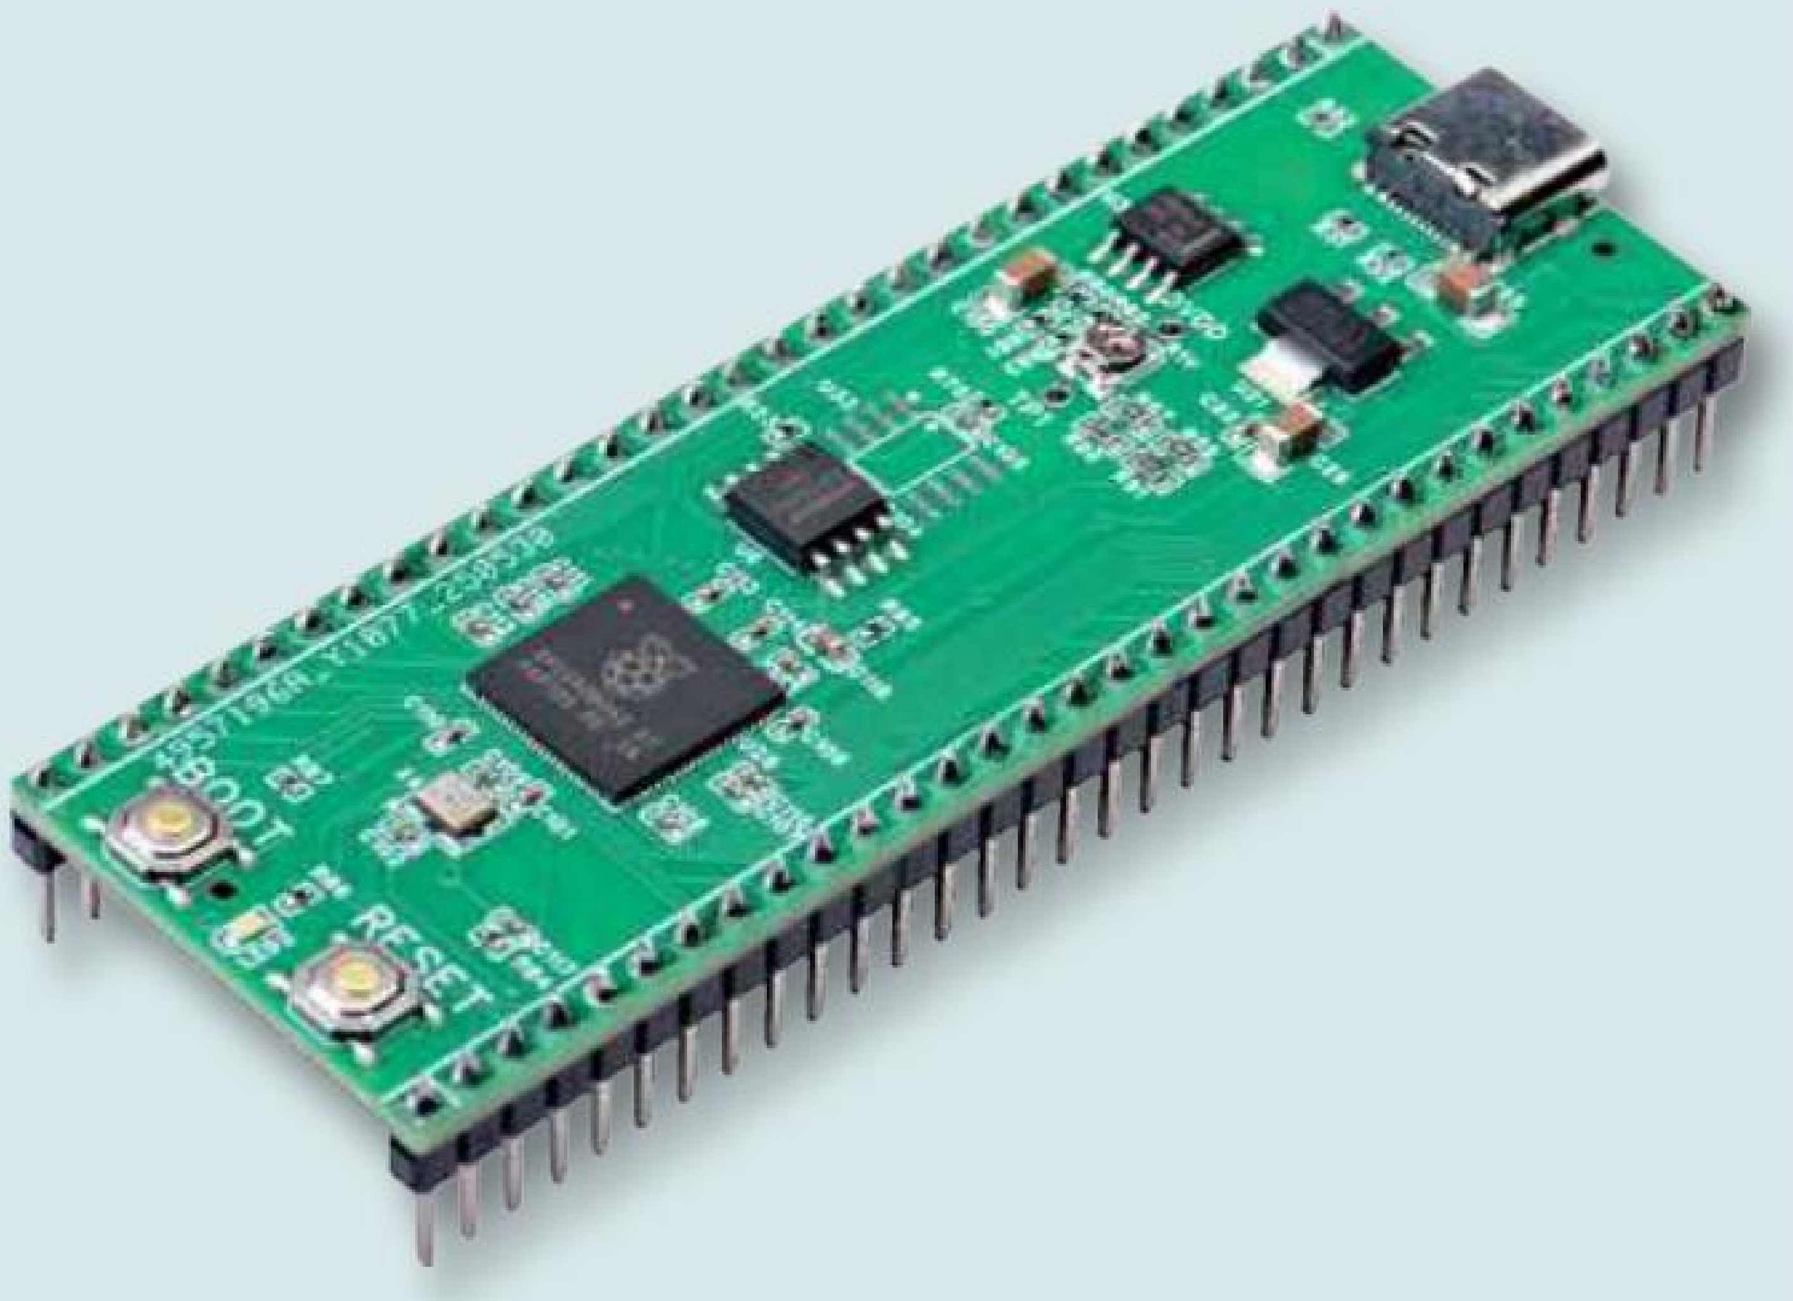

CR 2032
DURACELL

Projects and Circuits

RP2350B Development Board by Geoff Graham & Peter Mather 4
Think of this as the Raspberry Pi Pico 2's bigger sibling – more pins, more I/O, more speed, more storage and more memory. It's perfect for breadboarding, too. Like the Picos, it can be purchased pre-assembled.

USB Solar Charging by Julian Edgar 20
This simple electronic project allows you to charge and power all your USB devices from a low-cost solar system. It's inexpensive to put together, and once you've built it, charging your devices won't cost a cent!

The SmartProbe by Andrew Levido 23
The SmartProbe is an extremely handy little device for making voltage and continuity measurements. It won't replace your multimeter, but it is designed to be the first piece of test equipment you reach for when debugging or repairing a circuit.

Series, Features and Columns

Audio Out by Jake Rothman 10
All-pass filters, part four – new PCBs & symmetricisers. We start by describing an updated all-pass filter board plus two polarity inversion switches that can be used with it. We then go into a detailed history of symmetricisers used for radio broadcasting.

Circuit Surgery by Ian Bell 34
Measuring circuit frequency response using a PC sound card, part 10: differential amplifiers and line drivers. We continue looking at differential line drivers and also consider differential receivers, including practical implementations of both.

Techno Talk by Max the Magnificent 42
Tracks and tribulations – a ramble through the evolution of PCB layout, from tape and scalpels to overcaffeinated autorouters and today's AI-powered tools.

Mains Earthing Systems by Brandon Speedie 44
The Earth is an integral part of our power system. It can be used to improve electrical safety, reduce energy losses or save on the cost of a dedicated conductor. Here, we look at the different Earthing systems used worldwide and how they work.

Max's Cool Beans by Max The Magnificent 48
Weird & Wonderful Arduino Projects, part 19: marble mayhem and tilting tomfoolery. We add an accelerometer to our Retro Games Console so we can play by tipping it. This IMU can also measure the rate of rotation and even magnetic orientation.

Teach-In 12.9: the World of Wireless by Mike Tooley 56
RF sources and oscillators: this month, we take a deep dive into RF sources, explaining how to get the best out of a variety of phase-locked loop (PLL) signal and marker/generator modules, plus direct digital synthesis (DDS) chips.

The benefits of PCB inspection reports by Steve Mansfield-Devine (PCBWay) 66
However much you test your product designs, there will always be factors outside your control that affect your final product. One of the most important of these is the quality of the PCBs you receive from your chosen fabrication house.

Power Electronics, part four by Andrew Levido 69
AC-to-DC conversion using rectifiers – rectifier-type AC-DC converters predate the DC-DC converters we have been studying so far, by many decades. This month, we will look into these deceptively simple circuits.

The Fox Report by Barry Fox 78
Balcony solar is popular in Germany and it will soon be coming to the UK.

Regulars and Services

Editorial 2
A disappointing downgrade for the venerable NE5532

Reader services – Editorial and Advertising Departments 2

Subscribe to Practical Electronics and save money 3

Practical Electronics PCB Service 80-82
PCBs for Practical Electronics projects

Classified ads and Advertiser index 83

Next month! – highlights of our next issue of *Practical Electronics* 84

Renewals paid by PayPal, Visa and Mastercard are now handled via our website: **www.electronpublishing.com**

Made in Australia & the UK.

Written in Britain, Australia, the USA and Ireland.

Read everywhere.

© Silicon Chip Publications Pty Ltd 2026

Copyright in all drawings, photographs, articles, technical designs, software and intellectual property published in *Practical Electronics* is fully protected, and reproduction or imitation in whole or in part are expressly forbidden.

The August 2026 issue of *Practical Electronics* will be published on Tuesday, 7th of July, 2026 (see p84).

Practical Electronics

Volume 55. No. 07
July 2026
ISSN 2632 573X

Editorial

Editorial offices
Electron Publishing (Australia)
Tel +61 2 9939 3295
Email pe@pemag.au
Web www.electronpublishing.com

Address mail to:
Electron Publishing (Australia)
PO Box 194, Matraville NSW 2036
Australia

Advertising enquiries
+61 2 9939 3295 pe@pemag.au

Editor Nicholas Vinen
Publisher Nicholas Vinen
Digital subscriptions Stewart Kearn **Tel** 07918 614662
Online Editor Alan Winstanley
Web Systems Kris Thain
Production Bao Smith
Technical staff Tim Blythman, John Clarke

Print subscriptions
Practical Electronics Subscriptions
PO Box 6337
Bournemouth BH1 9EH
United Kingdom
Tel 01202 087631
Email pesubs@selectps.com

Technical enquiries
We regret that technical enquiries cannot be answered over the telephone. We are unable to offer any advice on the use, purchase, repair or modification of commercial equipment or the incorporation or modification of designs published in the magazine.

Questions about articles or projects should be sent to the editor by email: pe@pemag.au

Projects and circuits
All reasonable precautions are taken to ensure that the advice and data given to readers is reliable. We cannot, however, guarantee it and we cannot accept legal responsibility for it.

Some projects and circuits published in *Practical Electronics* employ voltages that can be lethal. Do not build, test, modify or fix any mains-powered equipment unless you fully understand the safety aspects involved and you use an RCD (GFCI) adaptor.

Component supplies
Silicon Chip Publications may offer kits or other parts for making our projects, but not in all cases. When kits are not available, readers will need to find and source parts themselves. We advise readers to check that all parts are still available before commencing any project in a back-dated issue.

Advertisements
Although the proprietors and staff of *Practical Electronics* take reasonable precautions to protect the interests of readers by ensuring as far as practicable that advertisements are bona fide, the magazine and its publishers cannot give any undertakings in respect of statements or claims made by advertisers, whether these advertisements are printed as part of the magazine, or in inserts. The Publishers regret that under no circumstances will the magazine accept liability for non-receipt of goods ordered, or for late delivery, or for faults in manufacture.

Transmitters/bugs/telephone equipment
We advise readers that certain items of radio transmitting and telephone equipment which may be advertised in our pages cannot be legally used in the UK. Readers should check the law before buying any transmitting or telephone equipment, as a fine, confiscation of equipment and/or imprisonment can result from illegal use or ownership. The laws vary from country to country; readers should check local laws.

A disappointing downgrade for the venerable NE5532

Nothing goes on forever, so it's never really a surprise when we find out that electronic parts we like to use are becoming obsolete or have become basically unavailable.

Sometimes it's very disappointing to find that there is no real replacement for those parts; often there is a pin-compatible substitute but it is inferior in certain ways. Still, there's usually a solution – perhaps a device in a different package (maybe an SMD) can fill the role. It may not be convenient, but that's life.

However, I was fairly shocked when Jacob Rothman informed me that the venerable NE5532 op amp was being silently replaced with an inferior version. You can read more about this in his Audio Out column this month, starting on page 15 under the heading "An unfortunate news flash".

Given that the NE5532 has been in production for around 50 years, it isn't at all surprising that it may be facing extinction. But this time things are a bit different. Rather than explaining that they aren't going to make it any more and providing a (possibly inferior) alternative, they have started making what is essentially a new part under the same name.

I don't think that's acceptable, especially given that there are plenty of circuits where this new NE5532 (per the Rev.K data sheet released by Texas Instruments in December 2025) would not work, despite the old version working fine. Even worse, there are circuits where the old NE5532 would have been quite happy and the new one will go up in a puff of smoke!

That's because they've lowered the maximum supply voltage from ±22V (±20V recommended maximum) to ±18V (±15V recommended maximum). Perhaps even more bizarre, the 'new NE5532' has NPN input transistors compared to PNP input transistors on the old version. So the input leakage current will have the opposite polarity! It also lacks input protection diodes, presumably making it less robust.

The thing is that there really isn't a replacement for the NE5532 in DIP. Of course, there are plenty of other good low-noise audio op amps, but they're mostly SMD-only these days, and many are quite expensive too. The NE5532 is practically the 'last man standing' as a high-performance, low-cost part in DIP.

There is one alternative: the Nisshinbo NJM5532D. While it is still available in quantity, it is no longer being made. It's also considerably more expensive. Stock up while you can. We have!

Nicholas Vinen, Electron Publishing (Australia)*
* a division of Silicon Chip Publications Pty Ltd.

Practical Electronics

UK readers SAVE £1 on every issue

SUBSCRIBE NOW!

01202 087631

Take out a one-year subscription and save more than £10 over the year.
Even better – save over £1 per issue if you subscribe for two years – a total saving of £26.

Overseas rates represent exceptional value

You also:
- **Avoid any cover price increase for the duration of your subscription**
- **Get your magazine delivered to your door each month**
- **Guarantee a copy, even if your local newsagent sells out**

Order by:
- **Phone or post with a cheque, postal order or credit card**
- **Or order online at: www.electronpublishing.com**

Practical Electronics

SUBSCRIPTION PRICES

Subscriptions for delivery direct to any address in:

UK: 6-months £33.99, 12-months £59.99, 24-months £114.99

Europe Airmail: 6-months £38.99, 12-months £69.99, 24-months £129.99

Rest Of The World Airmail: 6-months £46.99, 12-months £84.99, 24-months £164.99

Cheques (in **£ sterling only**) payable to ***Select Publisher Services Ltd*** and sent to: Practical Electronics Subscriptions, PO Box 6337, Bournemouth BH1 9EH, United Kingdom

Tel: 01202 087631
Email: pesubs@selectps.com

Also via our online shop at**: www.electronpublishing.com**

Subscriptions start with the next available issue. We accept MasterCard or Visa.
(For past issues see the Back Issues page.)

DIGITAL SUBSCRIPTIONS

Electronic subscriptions are available from £29.99 for 12 months, more details at: **www.electronpublishing.com**

SUBSCRIPTION ORDER FORM

☐ **6 Months: UK £33.99, Europe £38.99 (Airmail), Rest of the World £46.99 (Airmail)**
☐ **1 Year: UK £59.99, Europe £69.99 (Airmail), Rest of the World £84.99 (Airmail)**
☐ **2 Years: UK £114.99, Europe £129.99 (Airmail), Rest of the World £164.99 (Airmail)**

To: **Practical Electronics Subscriptions**
PO Box 6337
Bournemouth BH1 9EH
United Kingdom

Tel: 01202 087631
Email: pesubs@selectps.com

I enclose payment of £ ....
(cheque/PO in £ sterling only)
payable to ***Select Publisher Services Ltd***

☐ Please charge my Visa/Mastercard

My card number is: ....
Please print clearly, and check that you have the number correct

Card Security Code .... Valid From Date........................
(The last 3 digits on or just under the signature strip)

Card Ex. Date ....

Name ....

Address ....

Post code .... Tel. ....

Email ....

Subscriptions start with the next available issue.

RP2350B
DEVELOPMENT BOARD

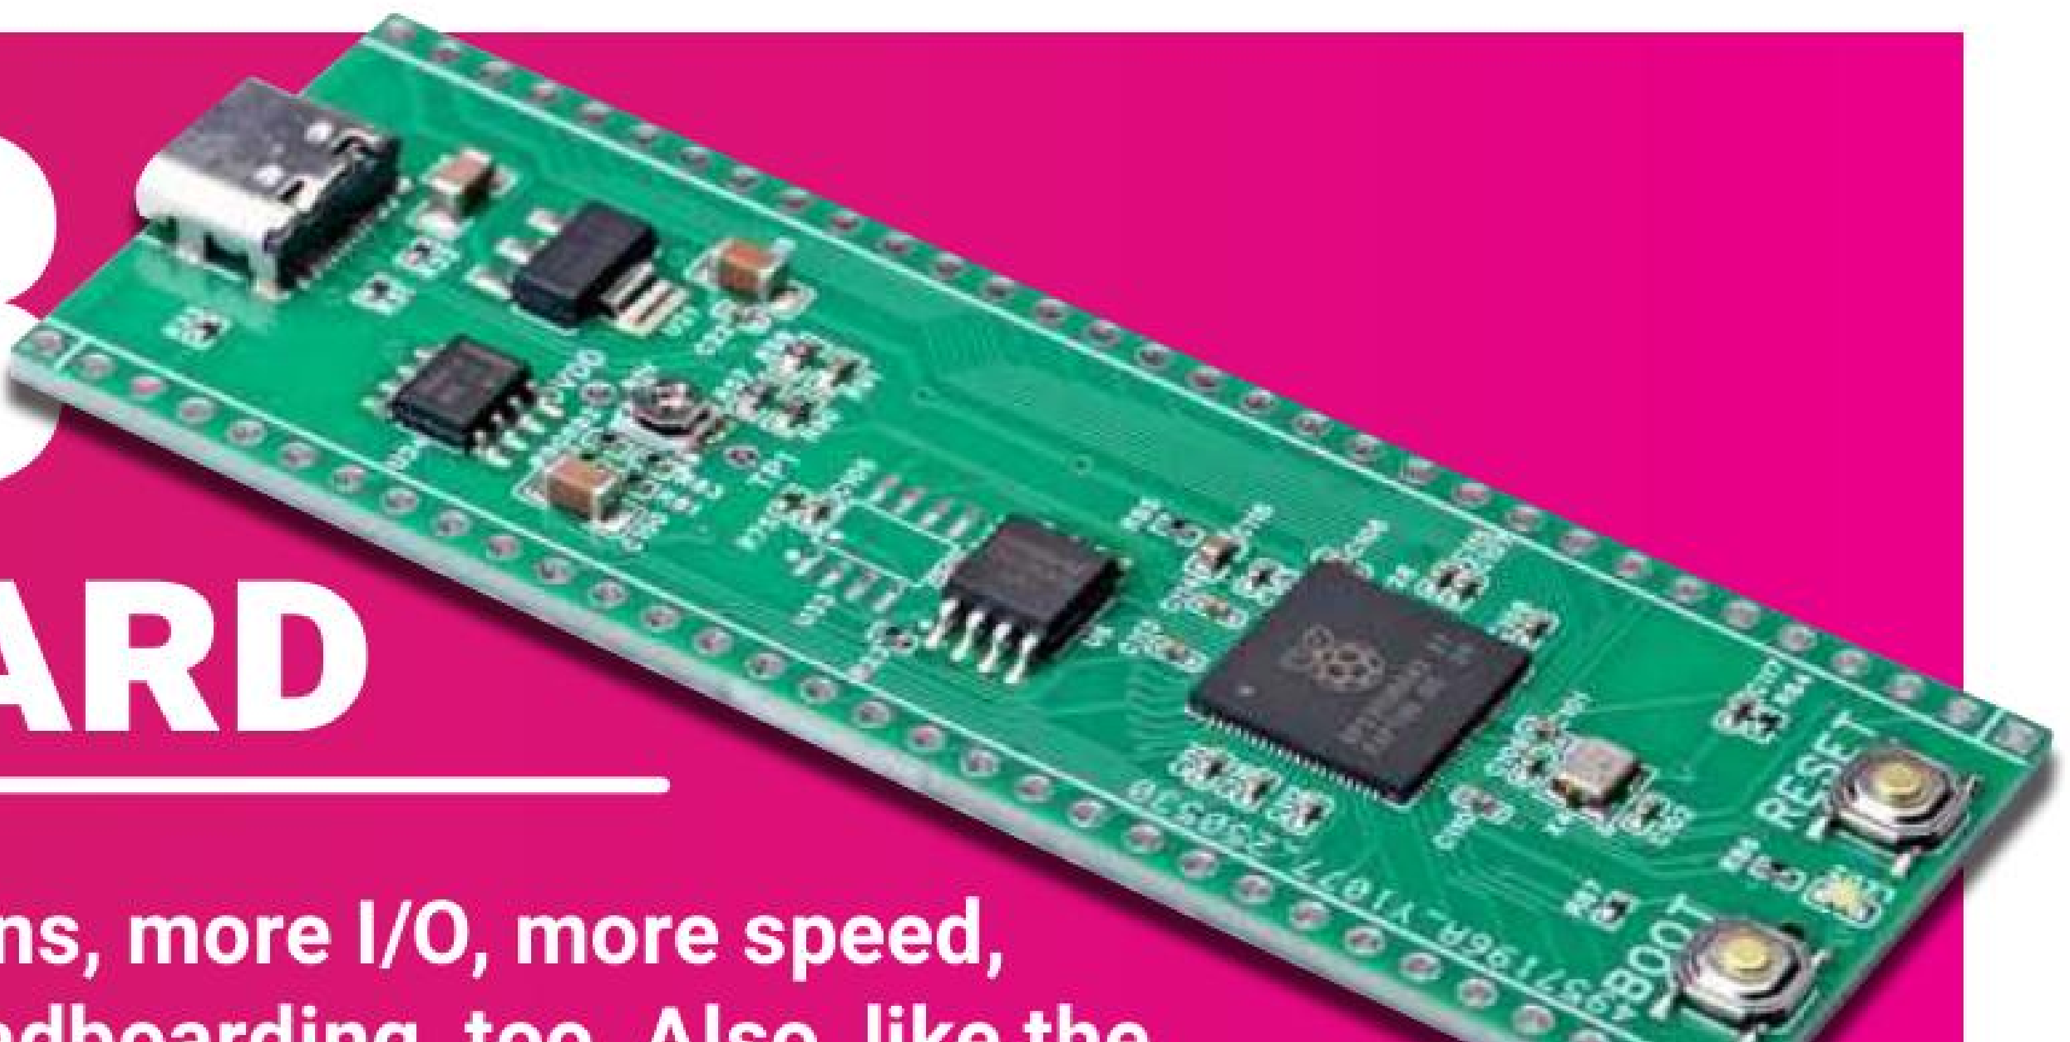

Think of this as the Pico 2's bigger sibling – more pins, more I/O, more speed, more storage and more memory. It's perfect for breadboarding, too. Also, like the Picos, it can be purchased pre-assembled!

The Raspberry Pi Pico is well known to our readers, and has been designed into many of our projects as a drop-in, self-contained compute module. It has been a runaway success for the Raspberry Pi Foundation, with over 4 million of this handy module sold since its launch in 2021.

In August 2024, they improved on it with the Pico 2, based on the new RP2350A processor, which has also been very popular.

The RP2350A processor also has a lesser-known sibling called the RP2350B, which is often overlooked. This has the same features as the A variant, but comes in a larger 80-pin QFN package with 48 I/O pins vs the 30 I/O pins of the RP2350A.

The Raspberry Pi Foundation does not offer a module based on this chip, but recently they have made it available for individual sale – so now we can design our own RP2350B based module. Thus, we present the RP2350B Development Board.

This is similar to the Raspberry Pi Pico 2, but it uses the RP2350B, with nearly all of its 48 I/O pins available for experimenters – a vast improvement over the Pico 2. This board can be used as a general-purpose compute module when you need a lot of I/O pins, or as a development board while designing a circuit around the RP2350B chip.

Features

The RP2350B Development Board is designed to suit solderless breadboards with two rows of 32 pins on a 2.54mm/0.1-inch pitch. This is similar to the standard Raspberry Pi Pico, and this layout can also be used as a plug-in module in your own PCB designs.

In our module, 47 of the RP2350B's 48 I/O pins are routed to the edge connectors. This is useful when you need a lot of I/O; for example, when driving a high-performance LCD panel with a parallel interface, or constructing your own multi-key keyboard. The one I/O pin not available, GPIO00, is used as the PSRAM IC chip select ($\overline{CS}$) signal.

On the edge pins of the module, we have also included numerous grounds, plus +3.3V and +5V outputs.

One extra benefit of the RP2350B is that eight of the I/O pins are capable of analog measurements (vs four on the RP2350A), and these are also available for use in your programs. The diagram opposite lists the full capabilities of each I/O pin.

The module is self-contained, including a 3.3V regulator. The input power supply is nominally 5V, but a range of 4.5-12V is acceptable. A USB-C socket is provided for power and loading the firmware. The firmware loading process works exactly the same as with the Raspberry Pi Pico, ie, you hold down the BOOT switch while plugging the USB into your desktop or laptop computer.

The flash memory used for storing programs and data in this design has a capacity of 16Mbytes (the Pico 2 has 4Mbytes). The PicoMite BASIC interpreter occupies just 2Mbytes, which leaves plenty of flash free to create an internal "disk drive" with a capacity of about 14Mbytes.

Our design also supports an 8Mbyte PSRAM chip. This sits on the same quad SPI bus as the flash memory, and can be used to add to the internal RAM of the RP2350B.

Overclocking

The RP2350B Development Board is designed for overclocking, which means running the processor cores at a higher clock speed than specified in the data sheet. This enables the module to be used in high-performance applications, such as generating DVI/HDMI video.

The default speed for the RP2350B is 150MHz, but some people have claimed to have overclocked it to over 600MHz. A more reasonable goal is the 372MHz needed to generate DVI/HDMI video.

To support the faster speeds, our design uses an adjustable linear regulator to generate the digital core supply voltage (DVDD). This powers the chip's core digital logic, and in our design, can be accurately adjusted from 1.1V to over 1.4V. Higher voltages allowing the CPU to run faster.

In the Pico 2, this voltage is provided by an on-chip switching regulator,

- **Processor: Raspberry Pi RP2350B**
- **Cores: two ARM Cortex-M33 and two Hazard3 RISC-V**
- **Clock Speed: default 150MHz; overclockable up to around 400MHz**
- **Flash memory: 16Mbytes**
- **RAM: 520kiB, expandable to over 8MiB**
- **I/O pins: 47 (eight with analog capability)**
- **I/O connectors: two rows of 32 pins, 2.54mm/0.1-inch pitch, 25.4mm/1-inch separation**
- **Power supply: 5V nominal (4.5-12V) @ 95mA (150MHz clock)**
- **Size: 82 × 28mm**

Words by Geoff Graham | Design by Peter Mather

RP2350B Assembled Board (£20): *a fully-assembled PCB with nearly everything from the parts list, except for the optional components.*

RP2350B Development Board

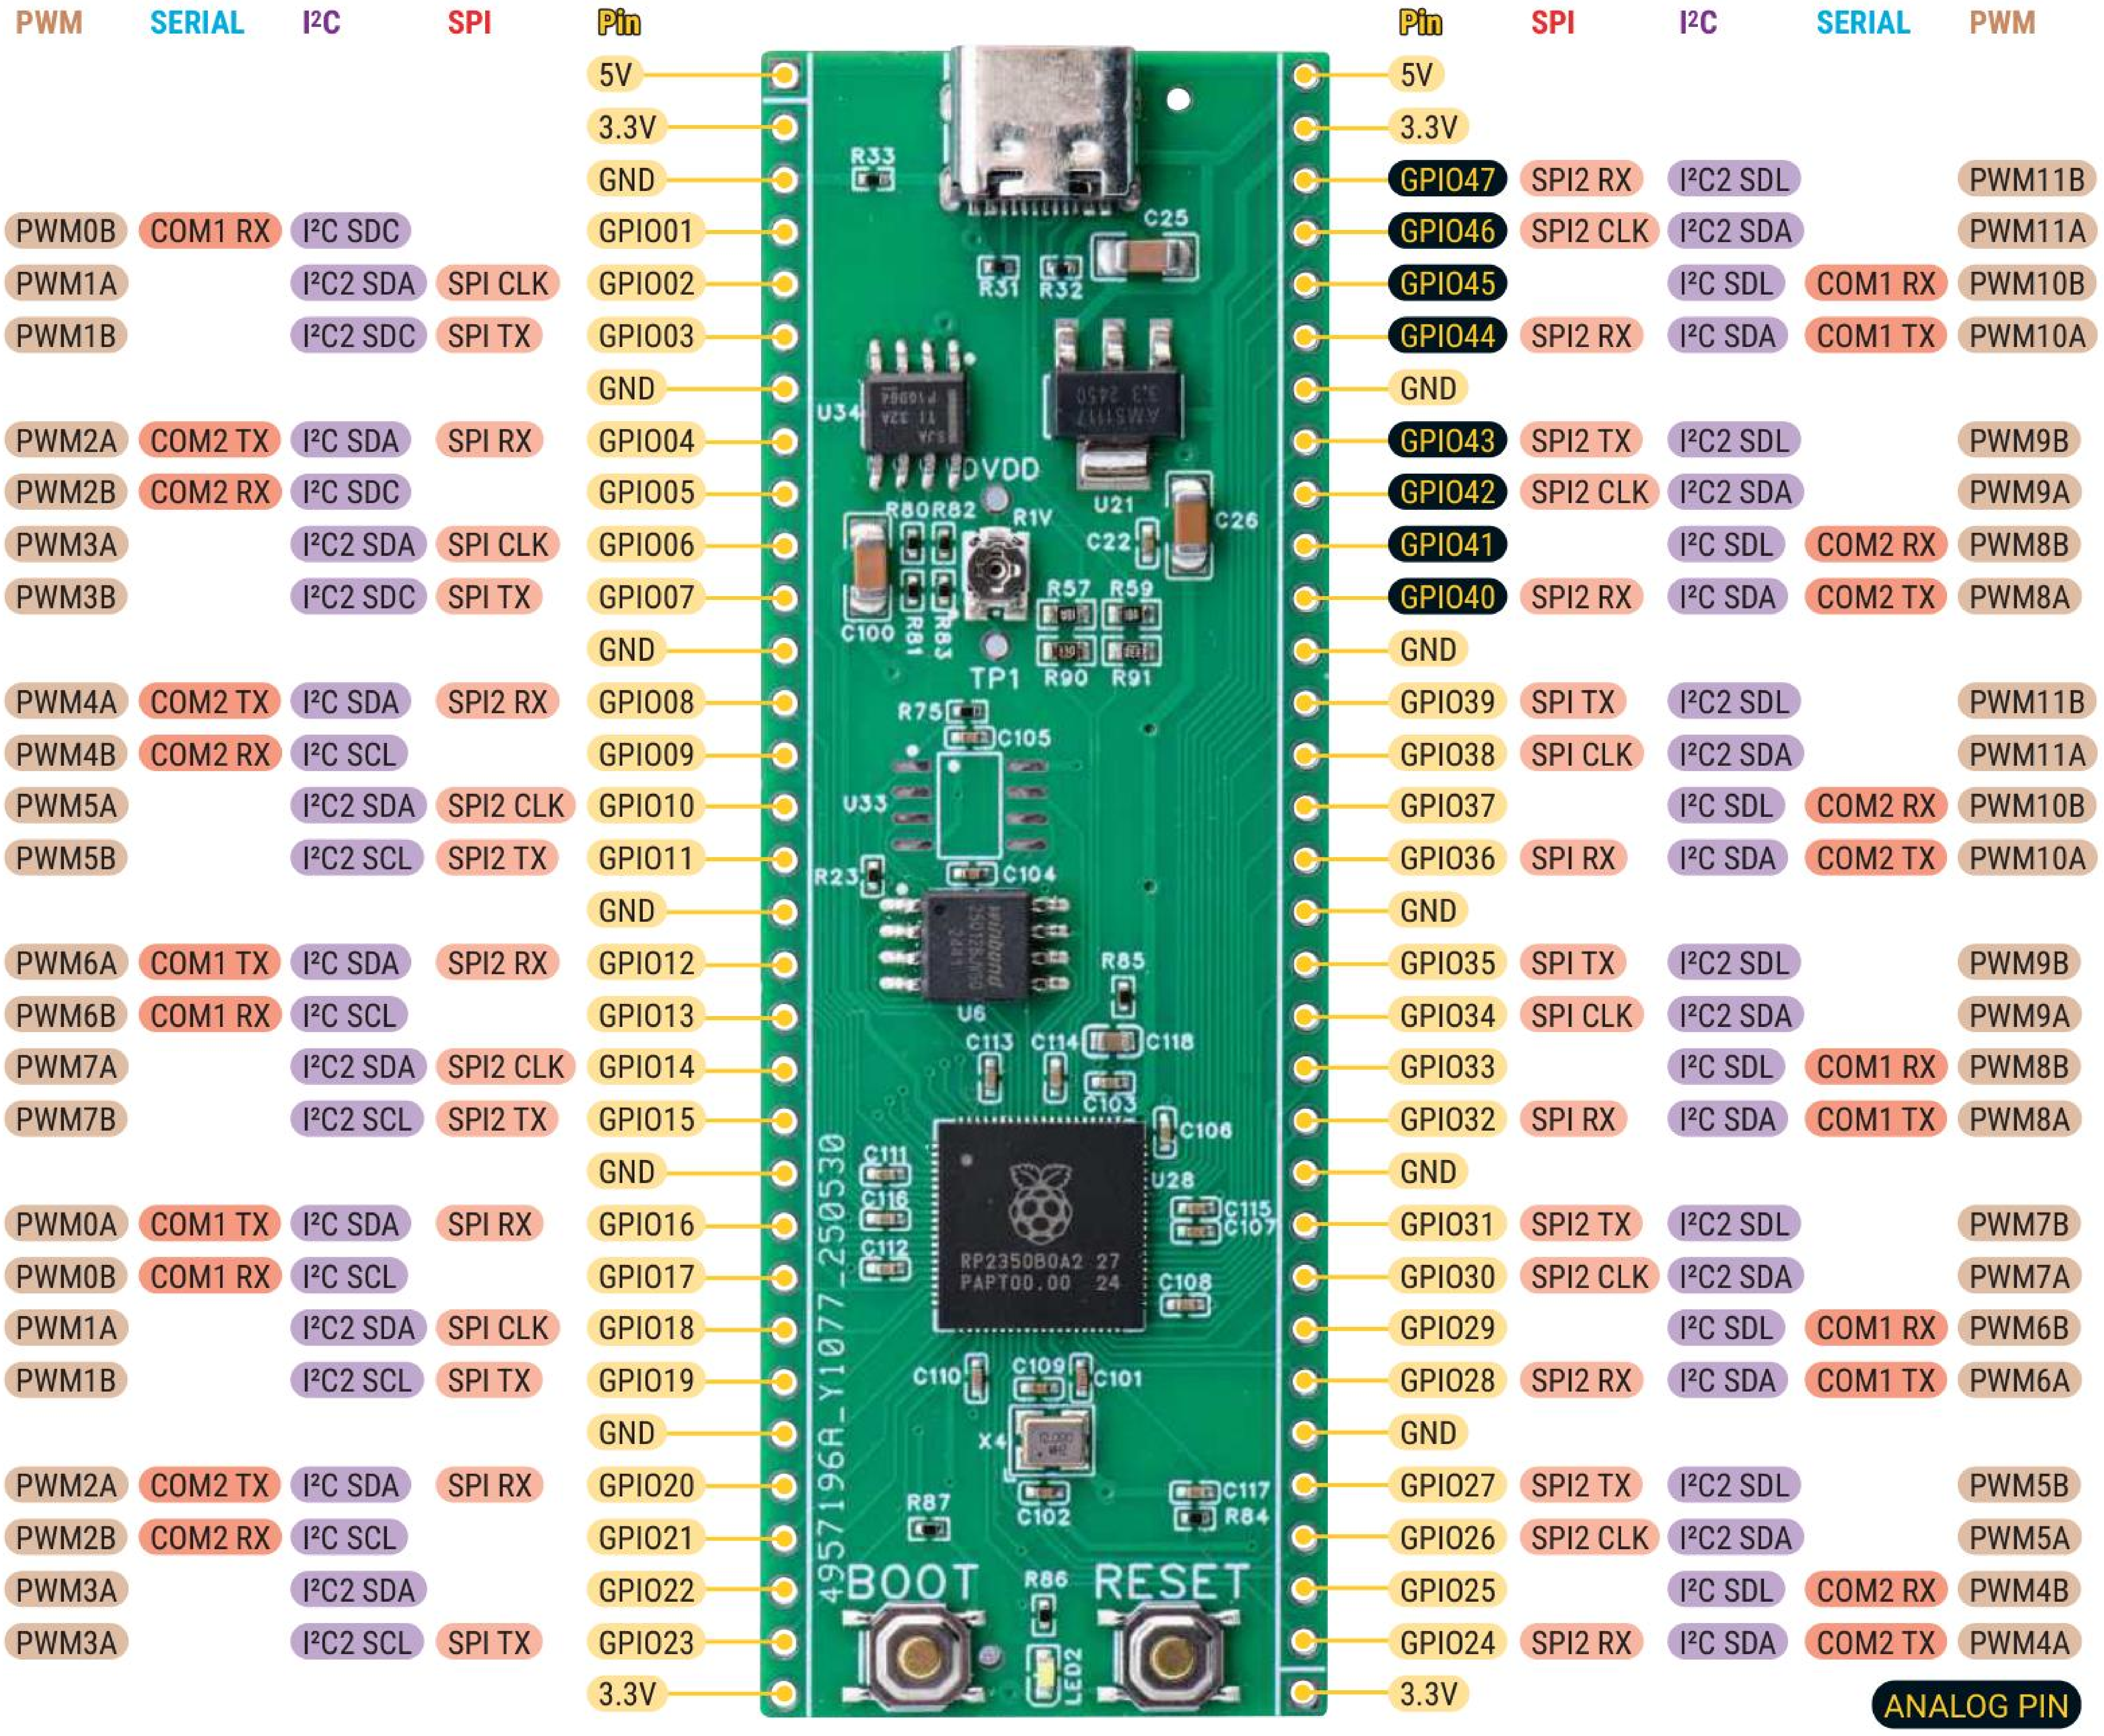
PWM SERIAL I^{2}C SPI Pin
5V
3.3V
GND
PWM0B COM1 RX I^{2}C SDC GPIO01
PWM1A I^{2}C2 SDA SPI CLK GPIO02
PWM1B I^{2}C2 SDC SPI TX GPIO03
GND
PWM2A COM2 TX I^{2}C SDA SPI RX GPIO04
PWM2B COM2 RX I^{2}C SDC GPIO05
PWM3A I^{2}C2 SDA SPI CLK GPIO06
PWM3B I^{2}C2 SDC SPI TX GPIO07
GND
PWM4A COM2 TX I^{2}C SDA SPI2 RX GPIO08
PWM4B COM2 RX I^{2}C SCL GPIO09
PWM5A I^{2}C2 SDA SPI2 CLK GPIO10
PWM5B I^{2}C2 SCL SPI2 TX GPIO11
GND
PWM6A COM1 TX I^{2}C SDA SPI2 RX GPIO12
PWM6B COM1 RX I^{2}C SCL GPIO13
PWM7A I^{2}C2 SDA SPI2 CLK GPIO14
PWM7B I^{2}C2 SCL SPI2 TX GPIO15
GND
PWM0A COM1 TX I^{2}C SDA SPI RX GPIO16
PWM0B COM1 RX I^{2}C SCL GPIO17
PWM1A I^{2}C2 SDA SPI CLK GPIO18
PWM1B I^{2}C2 SCL SPI TX GPIO19
GND
PWM2A COM2 TX I^{2}C SDA SPI RX GPIO20
PWM2B COM2 RX I^{2}C SCL GPIO21
PWM3A I^{2}C2 SDA GPIO22
PWM3A I^{2}C2 SCL SPI TX GPIO23
3.3V
Pin SPI I^{2}C SERIAL PWM
5V
3.3V
GPIO47 SPI2 RX I^{2}C2 SDL PWM11B
GPIO46 SPI2 CLK I^{2}C2 SDA PWM11A
GPIO45 I^{2}C SDL COM1 RX PWM10B
GPIO44 SPI2 RX I^{2}C SDA COM1 TX PWM10A
GND
GPIO43 SPI2 TX I^{2}C2 SDL PWM9B
GPIO42 SPI2 CLK I^{2}C2 SDA PWM9A
GPIO41 I^{2}C SDL COM2 RX PWM8B
GPIO40 SPI2 RX I^{2}C SDA COM2 TX PWM8A
GND
GPIO39 SPI TX I^{2}C2 SDL PWM11B
GPIO38 SPI CLK I^{2}C2 SDA PWM11A
GPIO37 I^{2}C SDL COM2 RX PWM10B
GPIO36 SPI RX I^{2}C SDA COM2 TX PWM10A
GND
GPIO35 SPI TX I^{2}C2 SDL PWM9B
GPIO34 SPI CLK I^{2}C2 SDA PWM9A
GPIO33 I^{2}C SDL COM1 RX PWM8B
GPIO32 SPI RX I^{2}C SDA COM1 TX PWM8A
GND
GPIO31 SPI2 TX I^{2}C2 SDL PWM7B
GPIO30 SPI2 CLK I^{2}C2 SDA PWM7A
GPIO29 I^{2}C SDL COM1 RX PWM6B
GPIO28 SPI2 RX I^{2}C SDA COM1 TX PWM6A
GND
GPIO27 SPI2 TX I^{2}C2 SDL PWM5B
GPIO26 SPI2 CLK I^{2}C2 SDA PWM5A
GPIO25 I^{2}C SDL COM2 RX PWM4B
GPIO24 SPI2 RX I^{2}C SDA COM2 TX PWM4A
3.3V
ANALOG PIN
BOOT
RESET

which is not suited to high clock speeds because of the electrical noise it generates. Additionally, it is difficult to implement, as it requires a specialised inductor, which is hard to find.

In our design, the DVDD voltage is provided by a TPS7A7002DDAR linear regulator (REG34) that is both inexpensive and does not generate any electrical noise. The onboard trimming resistor (VR1) is used to set the DVDD voltage. Note that it is important that this is set before power is applied to the board. If DVDD is accidentally set too high, it can damage the RP2350B chip.

We have also used an integrated crystal oscillator to generate the base clock of 12MHz for the RP2350B. This is different from Raspberry Pi Pico 2, which uses a simple crystal for this purpose. The integrated crystal oscillator provides a more stable clock with much less jitter. Jitter can be a problem when the base clock frequency is multiplied many times in the RP2350B to give the core CPU clock.

Development environments

All the familiar development environments used with the Raspberry Pi Pico 2 can be used with this board. This includes:

- The official Raspberry Pi C SDK for C/C++ development, which can be used from the command line on a desktop or laptop computer, or within popular integrated development environments like Visual Studio Code (VS Code), Eclipse and Clion.
- MicroPython, which is a full implementation of the Python 3 programming language running directly on the Development Board. This includes an interactive prompt to execute commands immediately via a USB serial port.
- Our own PicoMite firmware, which implements a feature-rich BASIC interpreter (MMBasic) with support for audio, LCD panels, SD cards, game controllers, HDMI/VGA video and PS2/USB keyboards. This firmware includes its own full-screen editor so programs can be developed, tested and run on the development board in a highly productive environment.

Circuit details

Fig.1 shows the full circuit for the RP2350B Development Module. At the centre is the RP2350B processor. All its general-purpose I/O pins are routed to the two connectors at the edges of the

The RP2350B has the same features as the RP2350A in the Pico 2 but has more pins, including 48 GPIOs.

RP2350B0A2 19
P6BB26.00 24

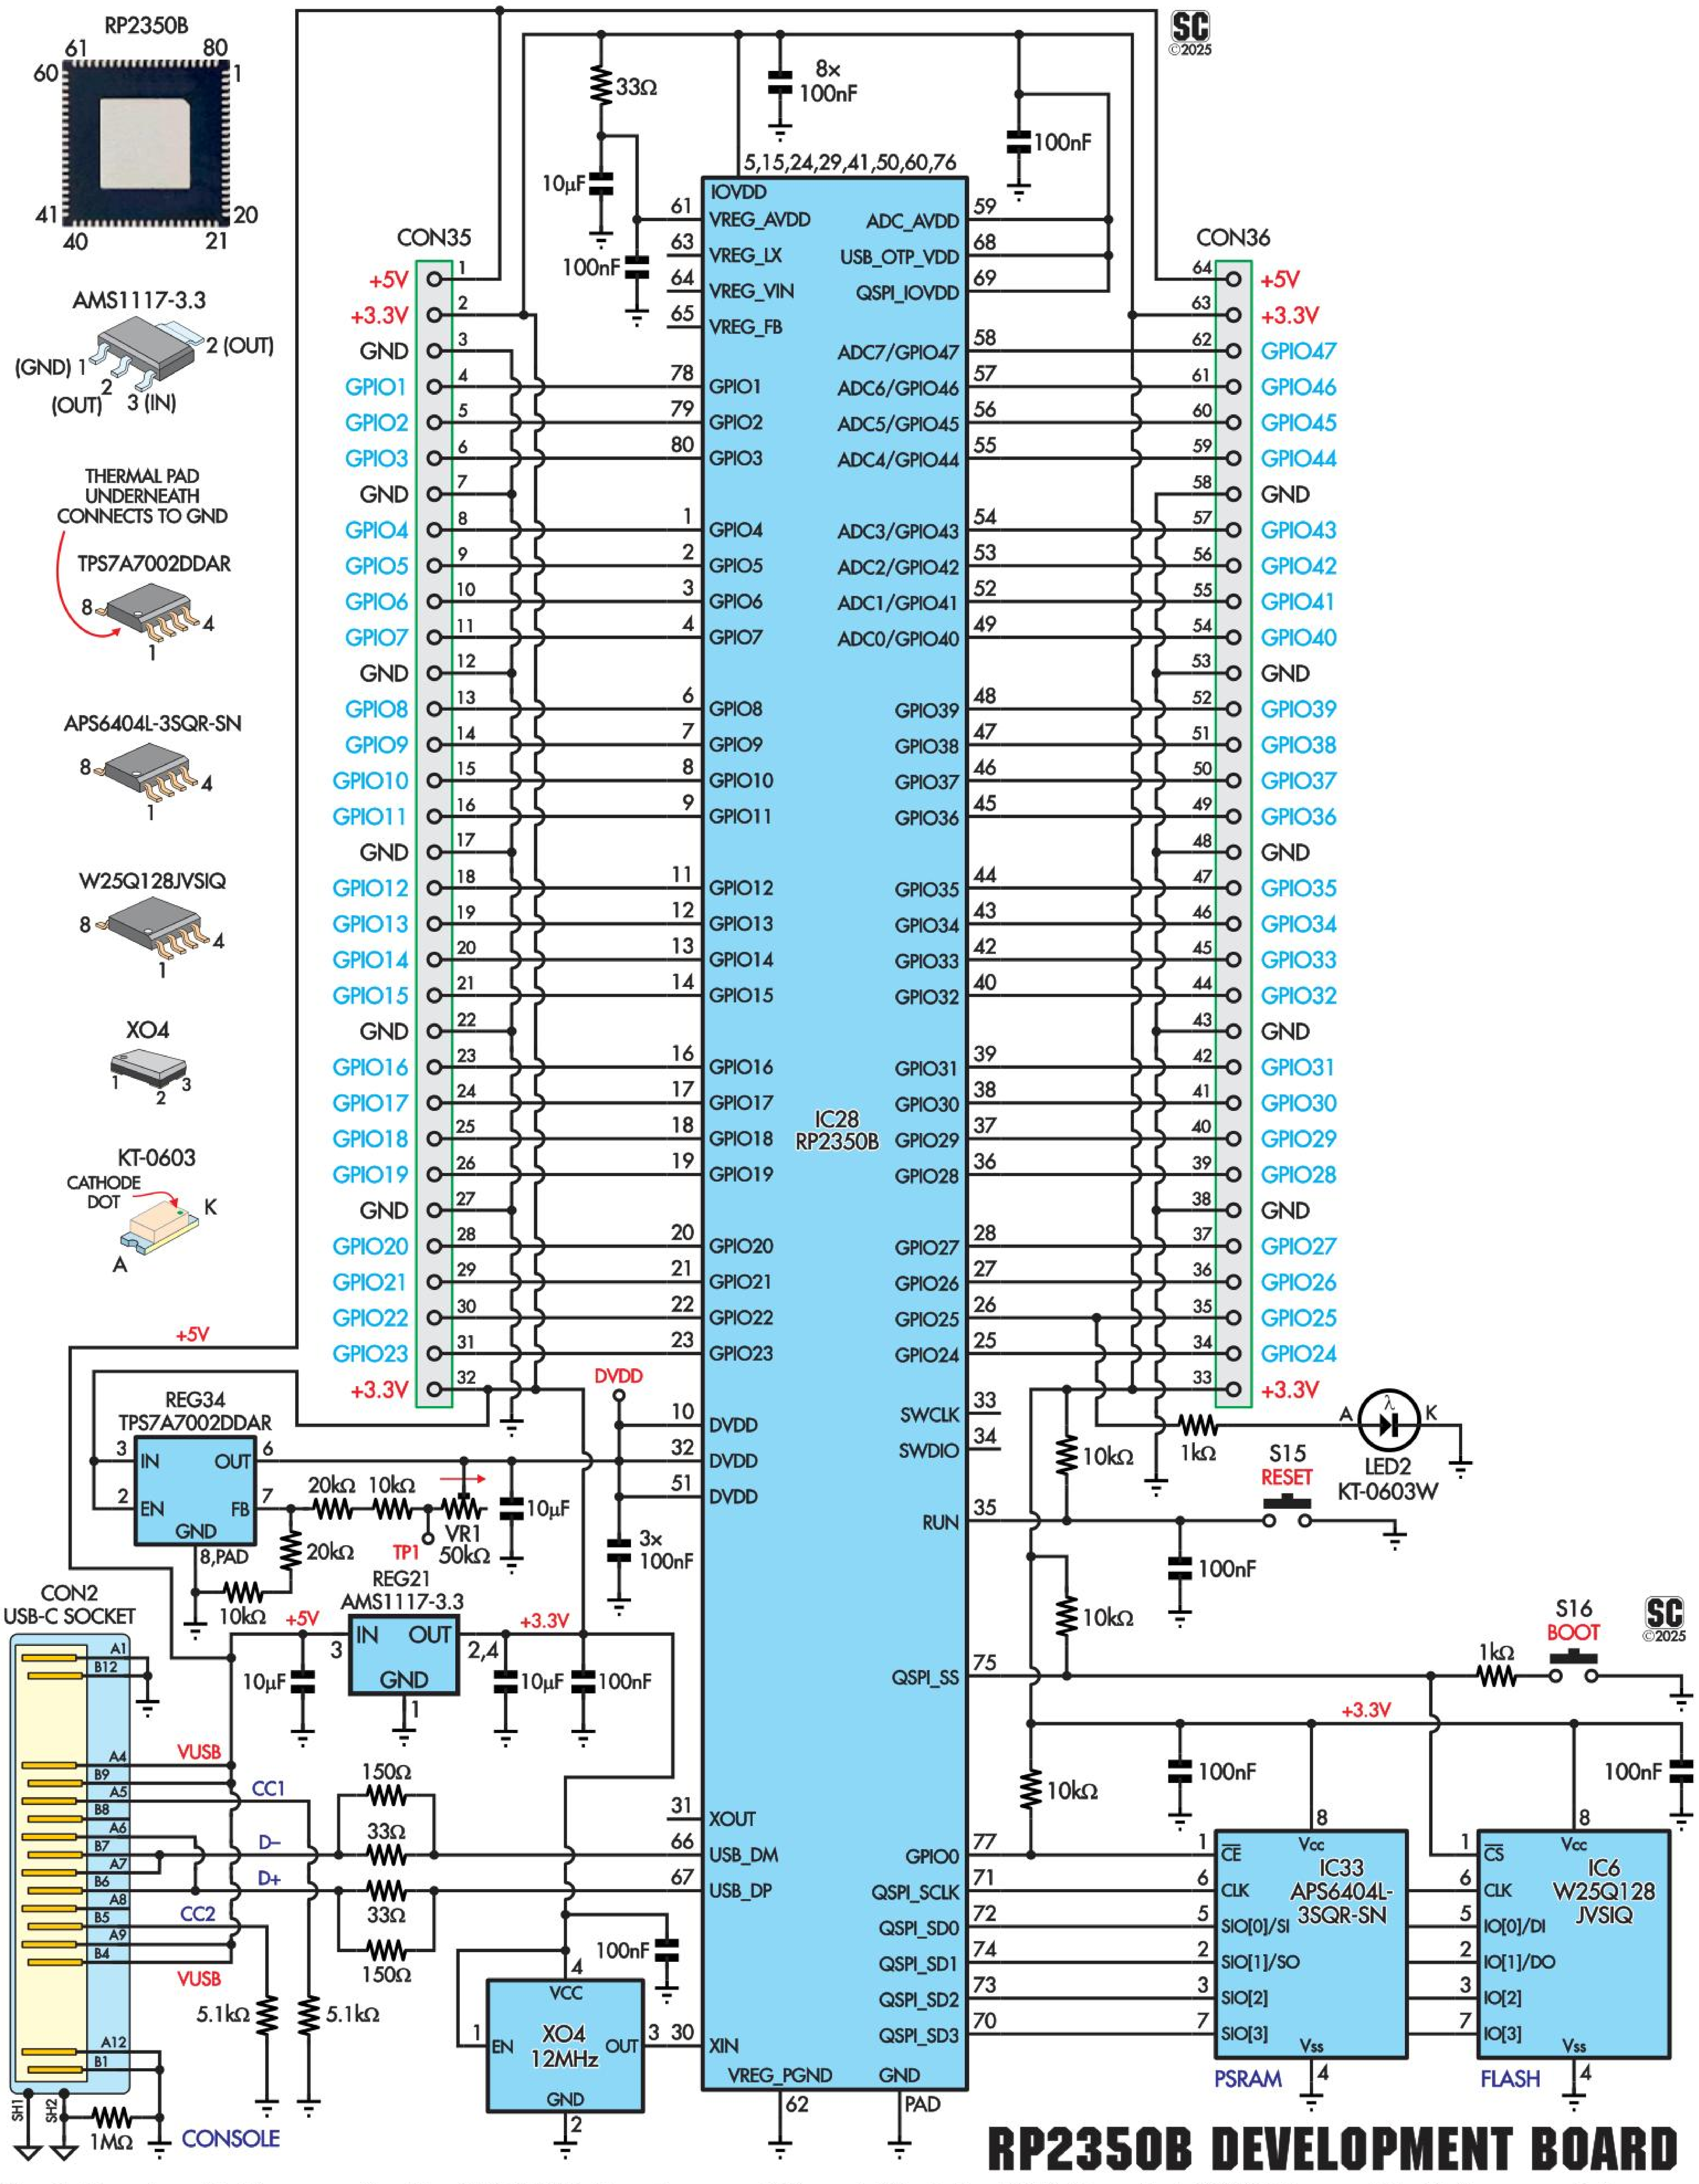
RP2350B
61
80
60
1
41
20
40
21
AMS1117-3.3
2 (OUT)
(GND) 1
2
(OUT)
3 (IN)
THERMAL PAD
UNDERNEATH
CONNECTS TO GND
TPS7A7002DDAR
8
4
1
APS6404L-3SQR-SN
8
4
1
W25Q128JVSIQ
8
4
1
XO4
1
2
3
KT-0603
CATHODE
DOT
K
A
SC
©2025
33Ω
8×
100nF
10µF
100nF
100nF
5,15,24,29,41,50,60,76
IOVDD
61 VREG_AVDD
63 VREG_LX
64 VREG_VIN
65 VREG_FB
ADC_AVDD 59
USB_OTP_VDD 68
QSPI_IOVDD 69
CON35
1 +5V
2 +3.3V
3 GND
4 GPIO1
5 GPIO2
6 GPIO3
7 GND
8 GPIO4
9 GPIO5
10 GPIO6
11 GPIO7
12 GND
13 GPIO8
14 GPIO9
15 GPIO10
16 GPIO11
17 GND
18 GPIO12
19 GPIO13
20 GPIO14
21 GPIO15
22 GND
23 GPIO16
24 GPIO17
25 GPIO18
26 GPIO19
27 GND
28 GPIO20
29 GPIO21
30 GPIO22
31 GPIO23
32 +3.3V
78 GPIO1
79 GPIO2
80 GPIO3
1 GPIO4
2 GPIO5
3 GPIO6
4 GPIO7
6 GPIO8
7 GPIO9
8 GPIO10
9 GPIO11
11 GPIO12
12 GPIO13
13 GPIO14
14 GPIO15
16 GPIO16
17 GPIO17
18 GPIO18
19 GPIO19
20 GPIO20
21 GPIO21
22 GPIO22
23 GPIO23
IC28
RP2350B
ADC7/GPIO47 58
ADC6/GPIO46 57
ADC5/GPIO45 56
ADC4/GPIO44 55
ADC3/GPIO43 54
ADC2/GPIO42 53
ADC1/GPIO41 52
ADC0/GPIO40 49
GPIO39 48
GPIO38 47
GPIO37 46
GPIO36 45
GPIO35 44
GPIO34 43
GPIO33 42
GPIO32 40
GPIO31 39
GPIO30 38
GPIO29 37
GPIO28 36
GPIO27 28
GPIO26 27
GPIO25 26
GPIO24 25
CON36
64 +5V
63 +3.3V
62 GPIO47
61 GPIO46
60 GPIO45
59 GPIO44
58 GND
57 GPIO43
56 GPIO42
55 GPIO41
54 GPIO40
53 GND
52 GPIO39
51 GPIO38
50 GPIO37
49 GPIO36
48 GND
47 GPIO35
46 GPIO34
45 GPIO33
44 GPIO32
43 GND
42 GPIO31
41 GPIO30
40 GPIO29
39 GPIO28
38 GND
37 GPIO27
36 GPIO26
35 GPIO25
34 GPIO24
33 +3.3V
+5V
REG34
TPS7A7002DDAR
3 IN
OUT 6
2 EN
FB 7
GND
8,PAD
20kΩ
10kΩ
VR1
50kΩ
TP1
20kΩ
10kΩ
10µF
DVDD
10 DVDD
32 DVDD
51 DVDD
3×
100nF
SWCLK 33
SWDIO 34
RUN 35
QSPI_SS 75
10kΩ
1kΩ
S15
RESET
A
K
LED2
KT-0603W
100nF
10kΩ
S16
BOOT
1kΩ
SC
©2025
+3.3V
100nF
100nF
10kΩ
CON2
USB-C SOCKET
A1
B12
A4
B9
A5
B8
A6
B7
A7
B6
A8
B5
A9
B4
A12
B1
SH1
SH2
1MΩ
CONSOLE
VUSB
CC1
D–
D+
CC2
VUSB
5.1kΩ
5.1kΩ
150Ω
33Ω
33Ω
150Ω
REG21
AMS1117-3.3
+5V
IN
OUT
+3.3V
3
2,4
GND
1
10µF
10µF
100nF
31 XOUT
66 USB_DM
67 USB_DP
3 30 XIN
VREG_PGND
62
GND
PAD
100nF
4
VCC
1 EN
XO4
12MHz
OUT
GND
2
GPIO0 77
QSPI_SCLK 71
QSPI_SD0 72
QSPI_SD1 74
QSPI_SD2 73
QSPI_SD3 70
1 CE
6 CLK
5 SIO[0]/SI
2 SIO[1]/SO
3 SIO[2]
7 SIO[3]
8 Vcc
Vss 4
IC33
APS6404L-
3SQR-SN
PSRAM
1 CS
6 CLK
5 IO[0]/DI
2 IO[1]/DO
3 IO[2]
7 IO[3]
8 Vcc
Vss 4
IC6
W25Q128
JVSIQ
FLASH
RP2350B DEVELOPMENT BOARD

Fig.1: the circuit diagram for the RP2350B Development Board/Module. USB-C socket CON2 is used both for supplying power (5V DC) and communicating with the RP2350B. The USB 5V supply is regulated to 3.3V by REG21 to power oscillator XO4, PSRAM IC33 and flash chip IC6. The RP2350B's nominally 1.1V core supply is generated by adjustable regulator REG34, so it can be increased for overclocking.

PCB (CON35/36), except for GPIO00, which is used as the chip select signal for the PSRAM (IC33).

One other pin (GPIO25) is special as it drives the onboard LED, LED2. If you do not need this indicating function, the I/O pin can be used as a general purpose I/O with the extra load presented by the LED and its 1kΩ current-limiting resistor.

The input power for the board is a nominal 5V DC; the RP2350B and the other chips on the board run from 3.3V, which is supplied by a simple AMS1117-3.3 low-dropout linear regulator.

The official Raspberry Pi Pico modules have a much more complex design for the power supply, using a switching regulator, but this can cause significant electrical noise that interferes with analog measurements and sensitive circuits such as audio input/output.

Most designs do not need the wide voltage range of the switching regulator, so our design avoids the noise problems and still provides a useful input supply voltage range of 4.5-12V.

The RP2350B needs another voltage supply called the digital core supply (DVDD), which we mentioned earlier as essential for overclocking. This powers the chip's core digital logic. In our design, it is provided by REG34, an adjustable linear regulator controlled by trimpot VR1.

IC6 is a W25Q128JVSIQ 128Mbit (16Mbyte) flash memory chip made by Winbond Electronics. It uses a quad SPI interface and is designed for true XIP (execute in place) operation, which allows the RP2350B to execute its program directly from this chip.

The W25Q128JVSIQ can operate with high clock speeds on the SPI interface (133MHz); this is important when overclocking the RP2350B. The RP2350B also has a built-in SRAM cache, which operates to mitigate the effect of the relatively slow quad SPI bus interface.

The BOOT pushbutton switch (S16) pulls the chip select line ($\overline{CS}$) low on the flash memory chip, which essentially disables the flash memory. When power is applied, the RP2350B will interpret the disabled memory as a signal to enter its bootloader mode, which is used to load a new firmware image.

IC33 is an optional external PSRAM chip (APS6404L-3SQR-SN), which sits on the same quad SPI bus as the flash memory chip. This has a capacity of 64Mbits (8Mbytes) and it can be used to augment the internal RAM of the RP2350B. How it is actually used depends on the running program.

Parts List – RP2350B Development Board

1 double-sided PCB coded 07107251, 82 × 28mm
2 momentary SMD tactile pushbutton switches (S15, S16)
  [XKB Connectivity TS-1187A-B-A-B]
1 USB-C USB 2.0 data + power socket (CON2) **[Kinghelm KH-TYPE-C-16P]**
1 50kΩ 3.8 × 3.6mm SMD trimpot (VR1) **[Bourns TC33X-2-503E]**
2 32-pin male headers, 2.54mm pitch (**optional**)

Semiconductors
1 128Mbit QSPI flash memory, SOIC-8 (IC6) **[Winbond W25Q128JVSIQ]**
1 Raspberry Pi RP2350B microcontroller, QFN-80 (IC28)
1 APS6404L-3SQR-SN 8MiB PSRAM, SOIC-8 (IC33; **optional**)
1 12MHz oscillator module, 3.2 × 2mm SMD-4 (XO4)
  [TOGNJING XOS32012000LT00351005]
1 AMS1117-3.3 low-dropout 3.3V linear regulator, SOT-223-3 (REG21)
1 TPS7A7002DDAR adjustable low-dropout voltage regulator, SOIC-8 (REG34)
1 white SMD LED, M1608/0603 size **[KT-0603W]**

Capacitors *(all SMD multilayer ceramic capacitors)*
3 10μF 50V X5R, M3216/1206 package **[Samsung CL31A106KBHNNNE]**
18 100nF 16V X7R, M1206/0402 package **[Samsung CL05B104KO5NNNC]**
1 10μF 10V X5R, M1608/0603 package
  [Samsung CL10A106KP8NNNC]

Resistors *(all SMD M1206/0402 ±1% unless noted)*
1 1MΩ
2 20kΩ
5 10kΩ
2 5.1kΩ
2 1kΩ
1 33Ω
2 33Ω (0603 size)
2 150Ω (0603 size)

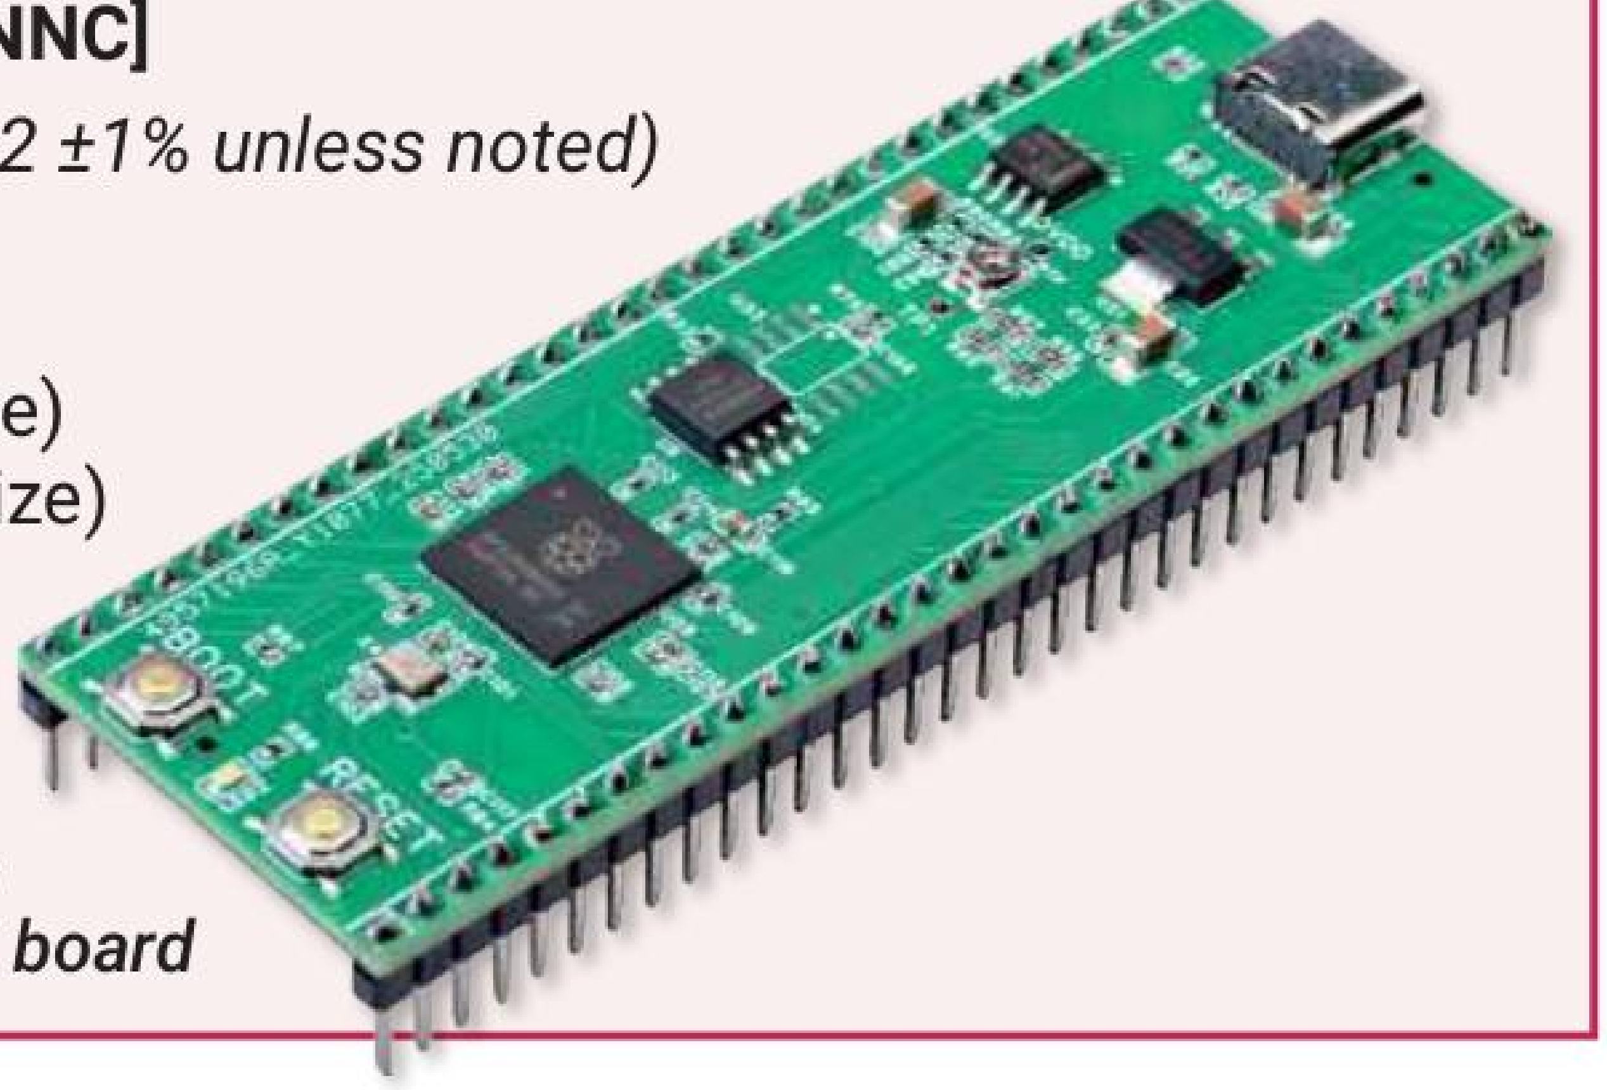

The finished RP2350B Development Board shown at 75% of actual size. The 32-pin headers are not included with the assembled board

For example, PicoMite BASIC will automatically add it to the general-purpose RAM seen by the BASIC interpreter, allowing for very large arrays to be defined.

Because the PSRAM must communicate over a serial interface, it is a lot slower than the internal RAM of the RP2350B. It also can limit the amount of overclocking that the board is capable of; however, it should still reach the speeds needed for generating DVI/HDMI video.

The internal RAM is normally more than enough for most applications, so for this and performance reasons, the PSRAM location is not populated in our design. However, it can be easily added to the BOM (Bill of Materials) for automated assembly or, as it is in an easy-to-solder 8-pin SOIC package, you can add it yourself.

This chip is available from Mouser for around $3 in small quantities: https://mouser.co.uk/ProductDetail/878-APS6404L-3SQR-SN

Building it

The RP2350B chip comes in an 80-pin QFP package, which is designed for automated surface-mount soldering. It can be hand soldered, but this is a challenge, even for someone skilled in SMD soldering.

So, practically speaking, you have two options for obtaining this module. You can purchase it fully assembled from the *Electron Publishing* website, or you can use an SMD assembly service such as JLCPCB's to build the board.

We recommend JLCPCB (https://jlcpcb.com) for the automated assembly, as they have proved reliable and reasonably priced in the past. JLCPCB also source the components at a good price and they do everything, including making the board, making the solder stencil, applying the solder paste,

Resistance (TP1-DVDD)	DVDD
6.0kΩ	1.10V
9.0kΩ	1.15V
12.0kΩ	1.20V
15.0kΩ	1.25V
18.0kΩ	1.30V
21.0kΩ	1.35V
24.0kΩ	1.40V
27.0kΩ	1.45V
30.0kΩ	1.50V
33.0kΩ	1.55V
36.0kΩ	1.60V

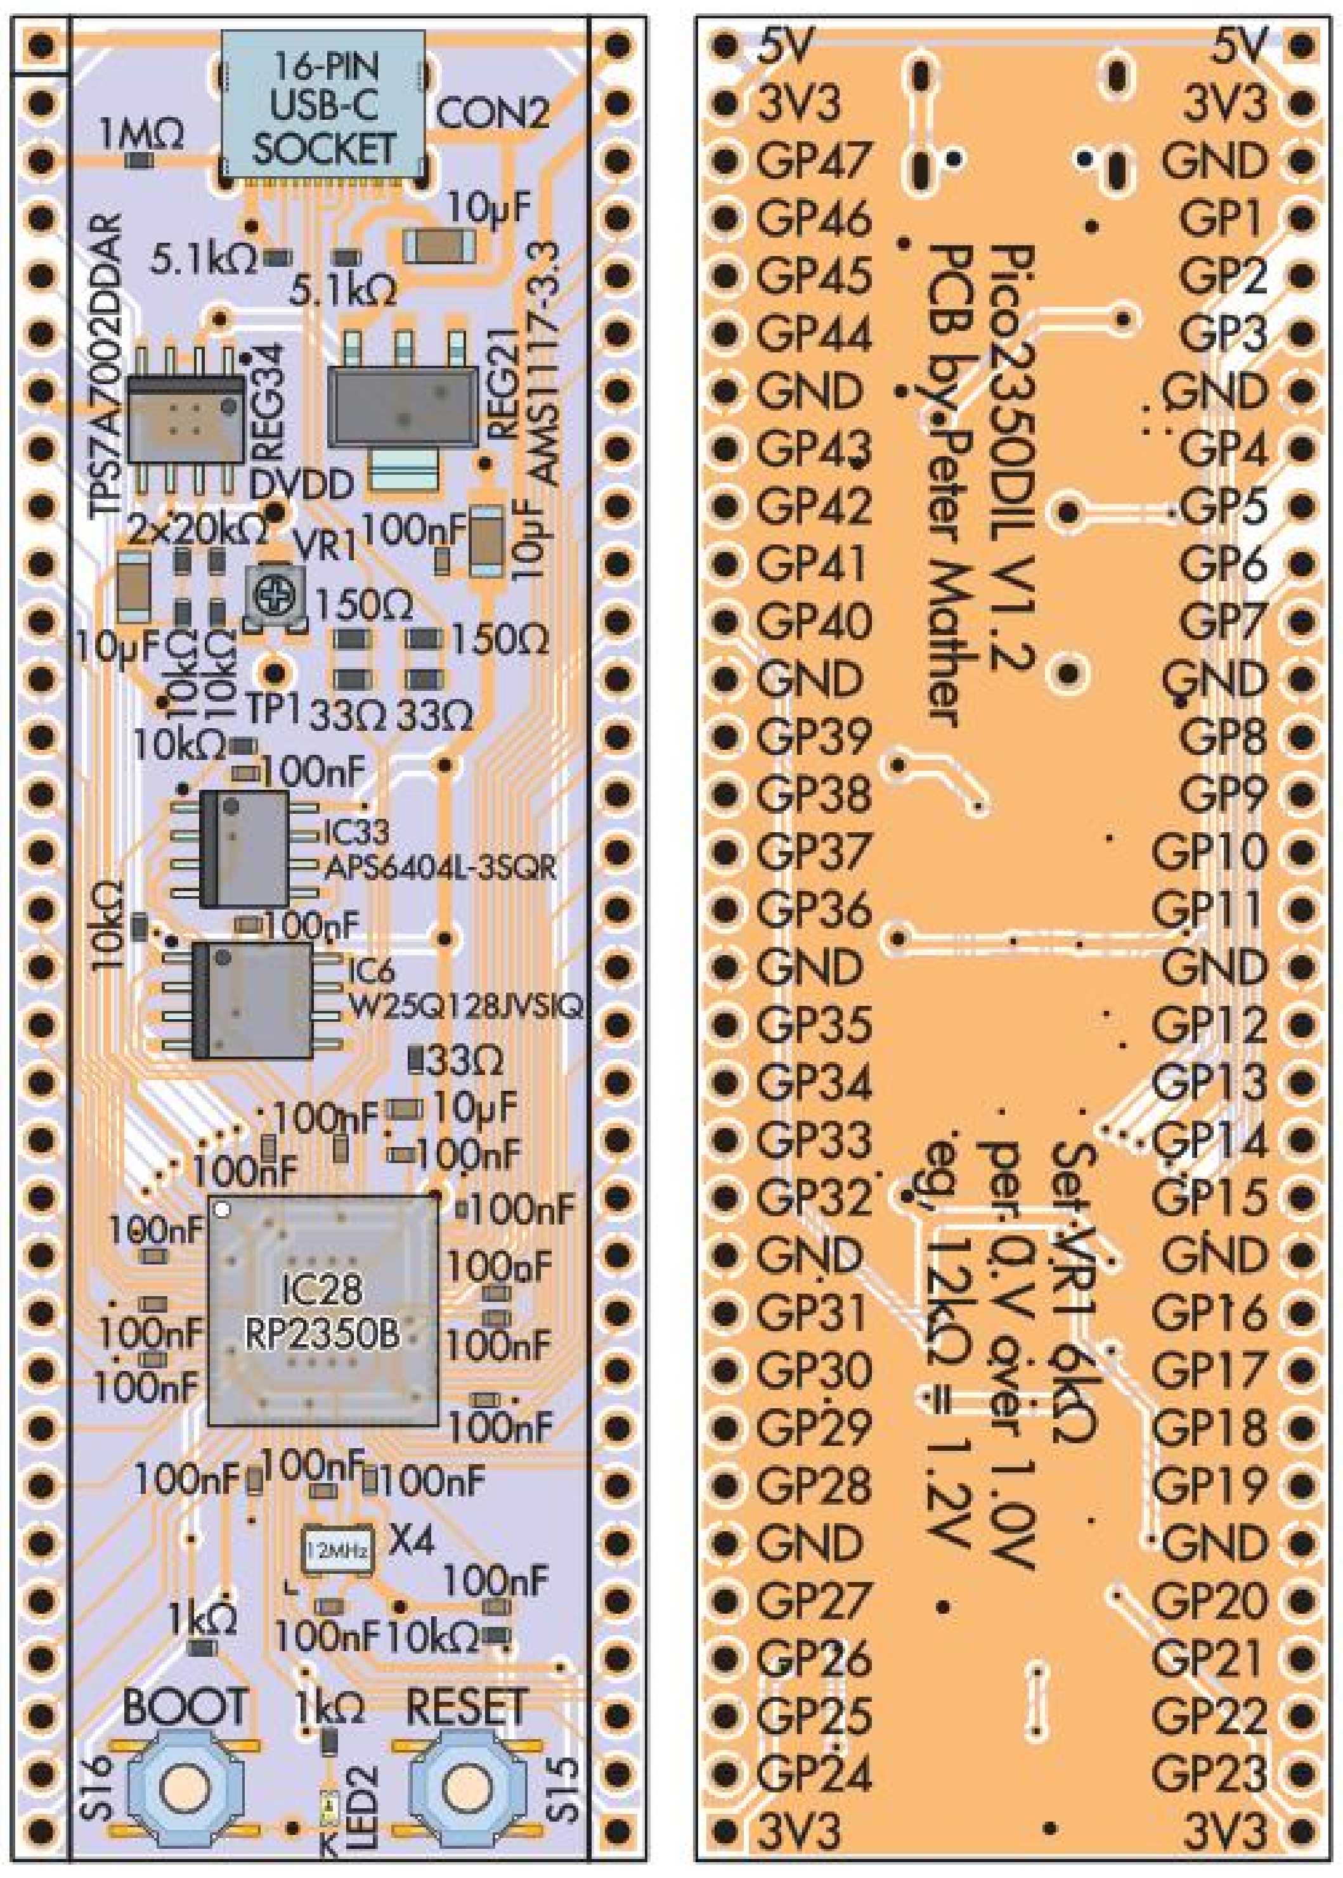
16-PIN USB-C SOCKET
CON2
1MΩ
5.1kΩ
5.1kΩ
10µF
TPS7A7002DDAR
REG34
REG21
AMS1117-3.3
DVDD
2×20kΩ
VR1
100nF
10µF
150Ω
150Ω
10µF
10kΩ
TP1
33Ω
33Ω
10kΩ
100nF
IC33
APS6404L-3SQR
10kΩ
100nF
IC6
W25Q128JVSIQ
33Ω
100nF
10µF
100nF
IC28
RP2350B
2MHz
X4
1kΩ
10kΩ
BOOT
1kΩ
RESET
S16
LED2
K
S15
5V
3V3
GP47
GP46
GP45
GP44
GND
GP43
GP42
GP41
GP40
GP39
GP38
GP37
GP36
GP35
GP34
GP33
GP32
GP31
GP30
GP29
GP28
GP27
GP26
GP25
GP24
GP1
GP2
GP3
GP4
GP5
GP6
GP7
GP8
GP9
GP10
GP11
GP12
GP13
GP14
GP15
GP16
GP17
GP18
GP19
GP20
GP21
GP22
GP23
Pico2350DIL V1.2
PCB by Peter Mather
Set VR1 6kΩ
per 0.1V over 1.0V
eg. 12kΩ = 1.2V

Fig.2 (above): the overlay diagram for the RP2350B Development Board. The table at the top of the page can be used as a reference for overclocking the RP2350B IC.

placing the components and reflow soldering them.

Their minimum quantity for assembly is two boards. However, if you do want a number of boards, it is hard to see why you would want to undertake the hand assembly of this board when the automated assembly option is relatively cheap.

To order boards from JLCPCB, you will need to download three files from siliconchip.au/Shop/10/2832 (they all come in one download package). The package includes a ZIP file with the Gerber files that contain the PCB design, a Bill of Materials spreadsheet listing all the parts required, and a CPL spreadsheet that contains the placement information for the pick-and-place robots.

Ordering assembled boards from the JLCPCB website is reasonably simple. Go to https://jlcpcb.com and create an account with them. Then, drag and drop the "RP2350B Development Board Gerbers.zip" file onto the "Add gerber file" box to the left of the Instant Quote button on their front page.

The website will process the file and display an image of the PCB. Click on the switch marked "PCB Assembly", then click "NEXT" until you reach a screen that prompts you to drag and drop the "RP2350B Development Board BOM.xlsx" and "RP2350B Development Board CPL.xlsx" files.

After supplying those files, you can continue and then accept all the defaults. However, you may wish to change the quantity of PCBs made (minimum of 5), and the number that you want assembled (minimum of 2).

Note that boards purchased from the *Electron Publishing* shop or assembled by JLCPCB will not include the two 32-pin headers required for use with a solderless breadboard. You will have to add these yourself (if required).

Adjusting the DVDD voltage

When you receive the boards, there is one adjustment that you need to make: using VR1 to set the digital core supply voltage (DVDD).

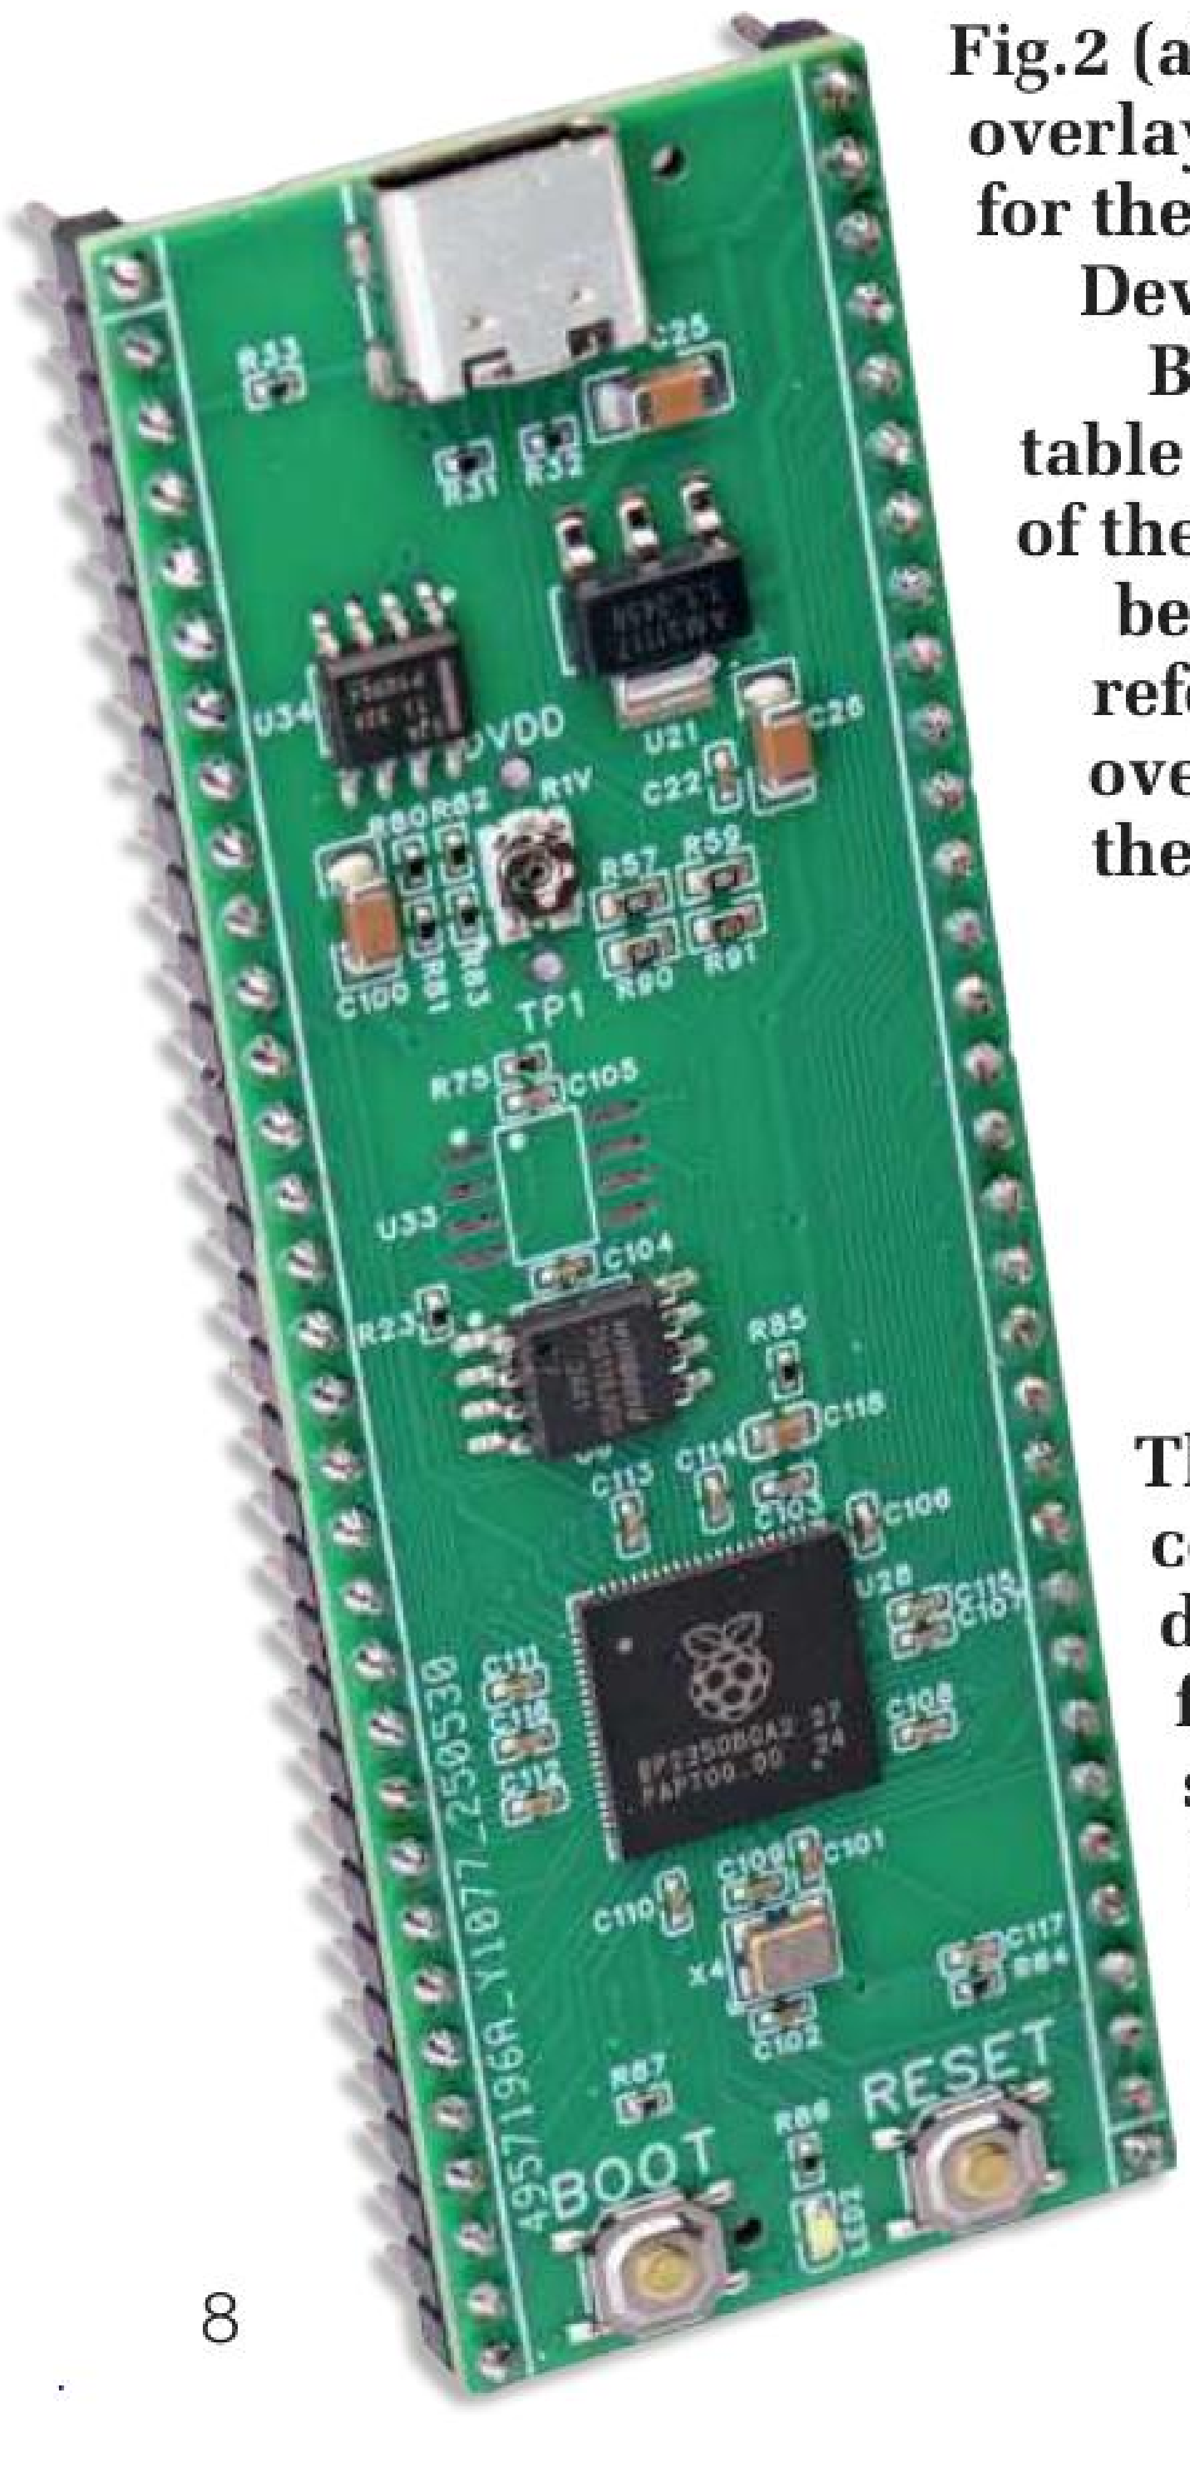
BOOT
RESET
TP1
DVDD

The RP2350B Development Board has all its components mounted on the top, with the pin designations listed on the reverse. It is available fully-assembled, apart from the two pin header strips. The module is quite small at 82×28mm (shown here enlarged for clarity). It is designed to suit solderless breadboards, with two rows of 32 pins on a 2.54mm/0.1in pitch, or it can be used as a plug-in module in your own PCB designs.

Important: the potentiometer must be adjusted before applying power to the board. Leaving it in a random position may damage the RP2350B chip.

Set your multimeter to the resistance mode and, with the board unpowered, place the leads across the test points marked DVDD and TP1. Adjust potentiometer VR1 to give a reading of 6kΩ. This will set DVDD to 1.1V, which is the standard voltage for a clock speed of 150MHz (the default for the RP2350B). This should also work for clock speeds up to 250MHz.

If you wish to overclock the RP2350B, you need to do two things: increase DVDD and set the desired clock speed in the program by setting CPU registers. Typically, a DVDD of 1.3V will allow the RP2350B to run up to 400MHz, with intermediate values suitable for clock speeds between 250MHz and 500MHz. The table at left lists some resistance values and the resultant DVDD voltages.

A maximum of 1.4V for DVDD should be safe. However, if you wish, you can try higher voltages with a risk of damaging the RP2350B processor. When the board is powered, this setting can be checked by measuring the voltage between the DVDD test point and any GND point.

Note that overclocking the RP2350B Development Board is not guaranteed, although all the samples we tested have reached a speed close to 400MHz. Also note that when a PSRAM chip is fitted, the maximum overclock speed will typically be slightly reduced.

Using the module

Power for the board is supplied via the USB-C connector. This is the normal mode when you are developing a program, as the board will be connected to a desktop or laptop computer that is used to load or edit the program.

When the board is used as an embedded controller (ie, not connected to a computer), power can be supplied via a 5V pin on the edge connector. This supply can be 4.5-12V. In this case, you cannot use the USB connector at the same time, as that would cause a conflict between the two power supplies and possibly damage your computer.

To prevent this possibility, you can use a switch or jumper to isolate your power supply whenever the USB connector is used.

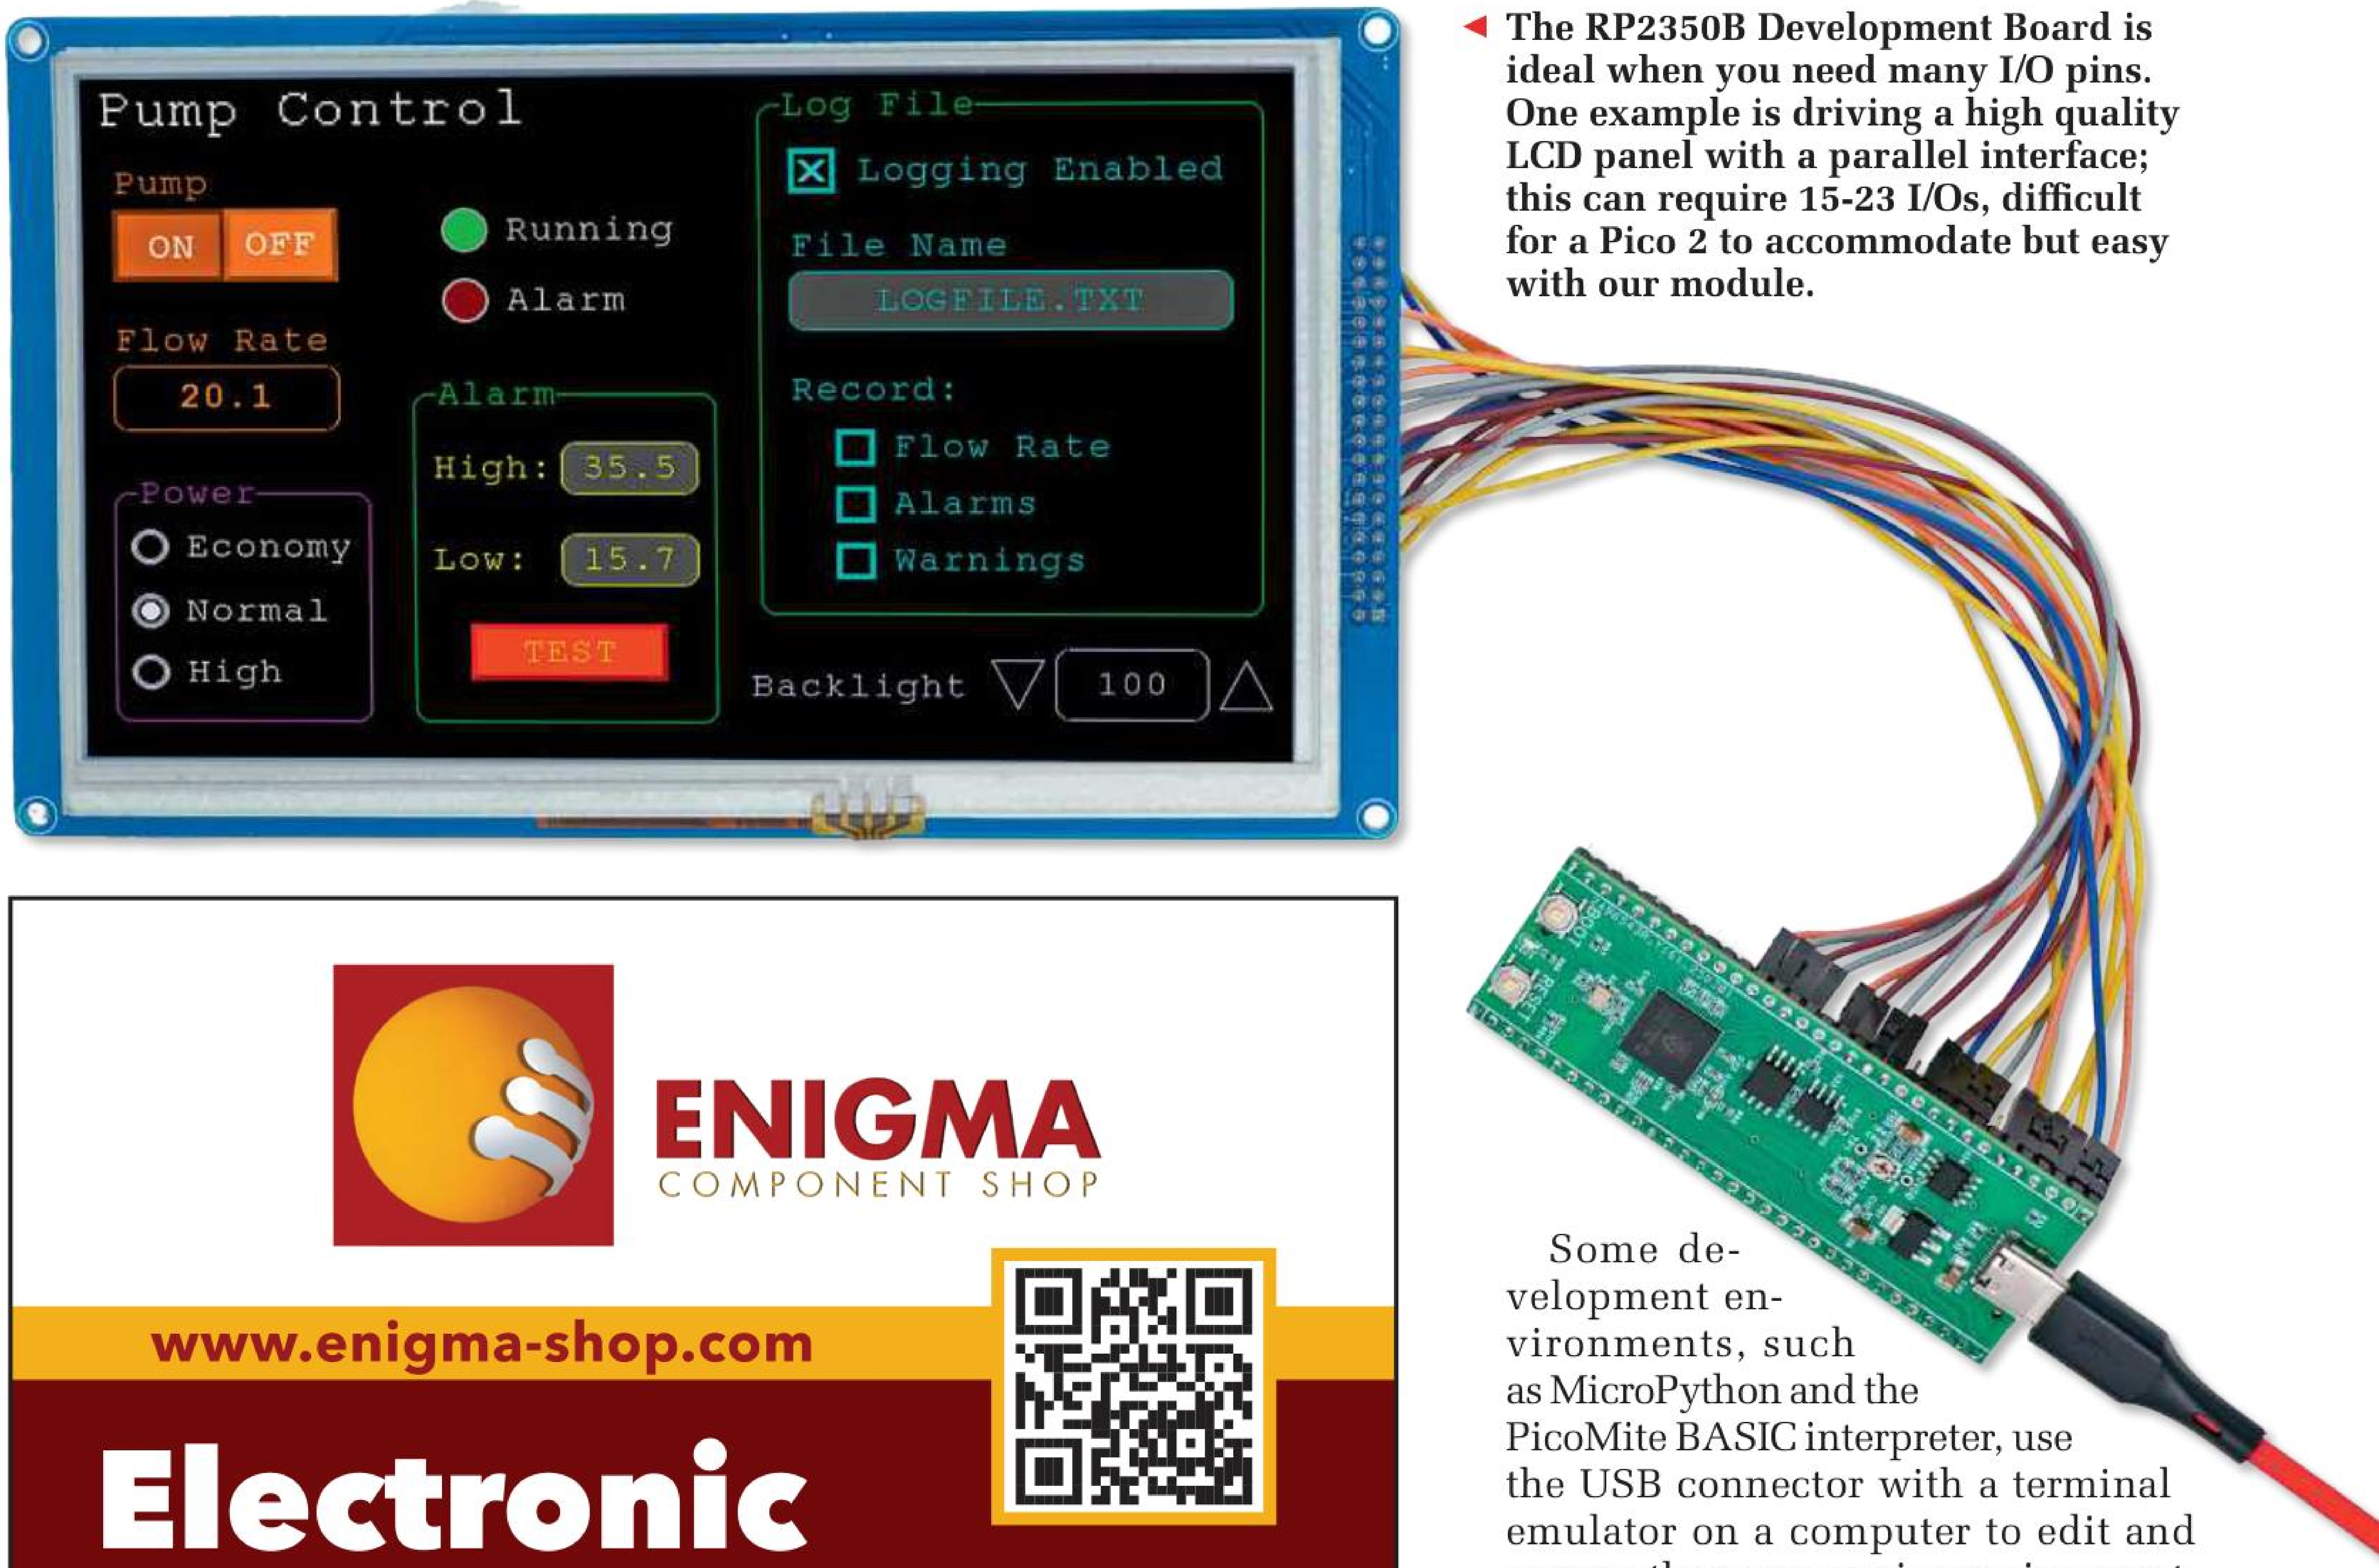
Pump Control
Pump
ON
OFF
Running
Alarm
Flow Rate
20.1
Alarm
High:
35.5
Low:
15.7
TEST
Power
○ Economy
● Normal
○ High
Log File
☒ Logging Enabled
File Name
LOGFILE.TXT
Record:
☐ Flow Rate
☐ Alarms
☐ Warnings
Backlight
100

The RP2350B Development Board is ideal when you need many I/O pins. One example is driving a high quality LCD panel with a parallel interface; this can require 15-23 I/Os, difficult for a Pico 2 to accommodate but easy with our module.

ENIGMA
COMPONENT SHOP
www.enigma-shop.com

Electronic Components for the Hobbyist!

Cable & Wire • Capacitors • Connectors • Filters
Kits & Modules • Hardware & Cooling
and much more...

READY TO BE REWARDED?
Join our loyalty scheme and **earn points** on every purchase you make which can be **used for discounts** in the future.

Read our reviews on Trustpilot

Some development environments, such as MicroPython and the PicoMite BASIC interpreter, use the USB connector with a terminal emulator on a computer to edit and manage the programming environment.

Other, hosted development environments, such as the C/C++ compiler, will build the program on the desktop or laptop computer, which then needs to be transferred to the module.

To load this firmware, you simply hold down the BOOT button while restarting the module (the RESET button is good for that) and then copy the firmware to the pseudo USB drive that is created on your computer by the RP2350B chip.

How you use the RP2350B Development Board will depend on the firmware that you have running on it. The PicoMite BASIC interpreter can be downloaded from siliconchip.au/Shop/6/833

This is a complete OS with a Microsoft BASIC compatible interpreter and extensive hardware support, including HDMI/VGA video, PS/2 and USB keyboards, touch-sensitive LCD panels, SD cards and much more.

For a full description of the PicoMite firmware and what it can do, read the article on the PicoMite 2 in the February 2026 issue, or visit the author's website at https://geoffg.net PE

AUDIO OUT

By Jake Rothman

All-pass filters, part four – new PCBs & symmetricisers

To make full use of the all-pass filter board presented last month, we need a better polarity-inversion switch. Coupling capacitors and pull-down resistors to stop offset voltages building up have been added to a switch mounting PCB, preventing switch clicks.

This means that the resistor, formerly R41 in Fig.34, described last month as being mounted under the board is no longer needed. It is now R4 on the switch board. This added switch circuit is shown in Fig.35, the overlay in Fig.36, and the final construction in Photo 26. It plugs into the new main PCB shown in Photo 27. These additional components have their own numbering.

This assembly is useful in its own right for adding audio change-over switching functions to any directly coupled circuitry.

There's provision for a dual-colour LED. When the switch is pushed in, it glows red and the polarity is reversed, while green's for normal. This LED is a two-terminal back-to-back red/green type, such as the Tayda A-1076. This is driven at around 2mA via a 10kΩ 0.6W current-limiting resistor R9, which gives enough brightness with modern panel LEDs.

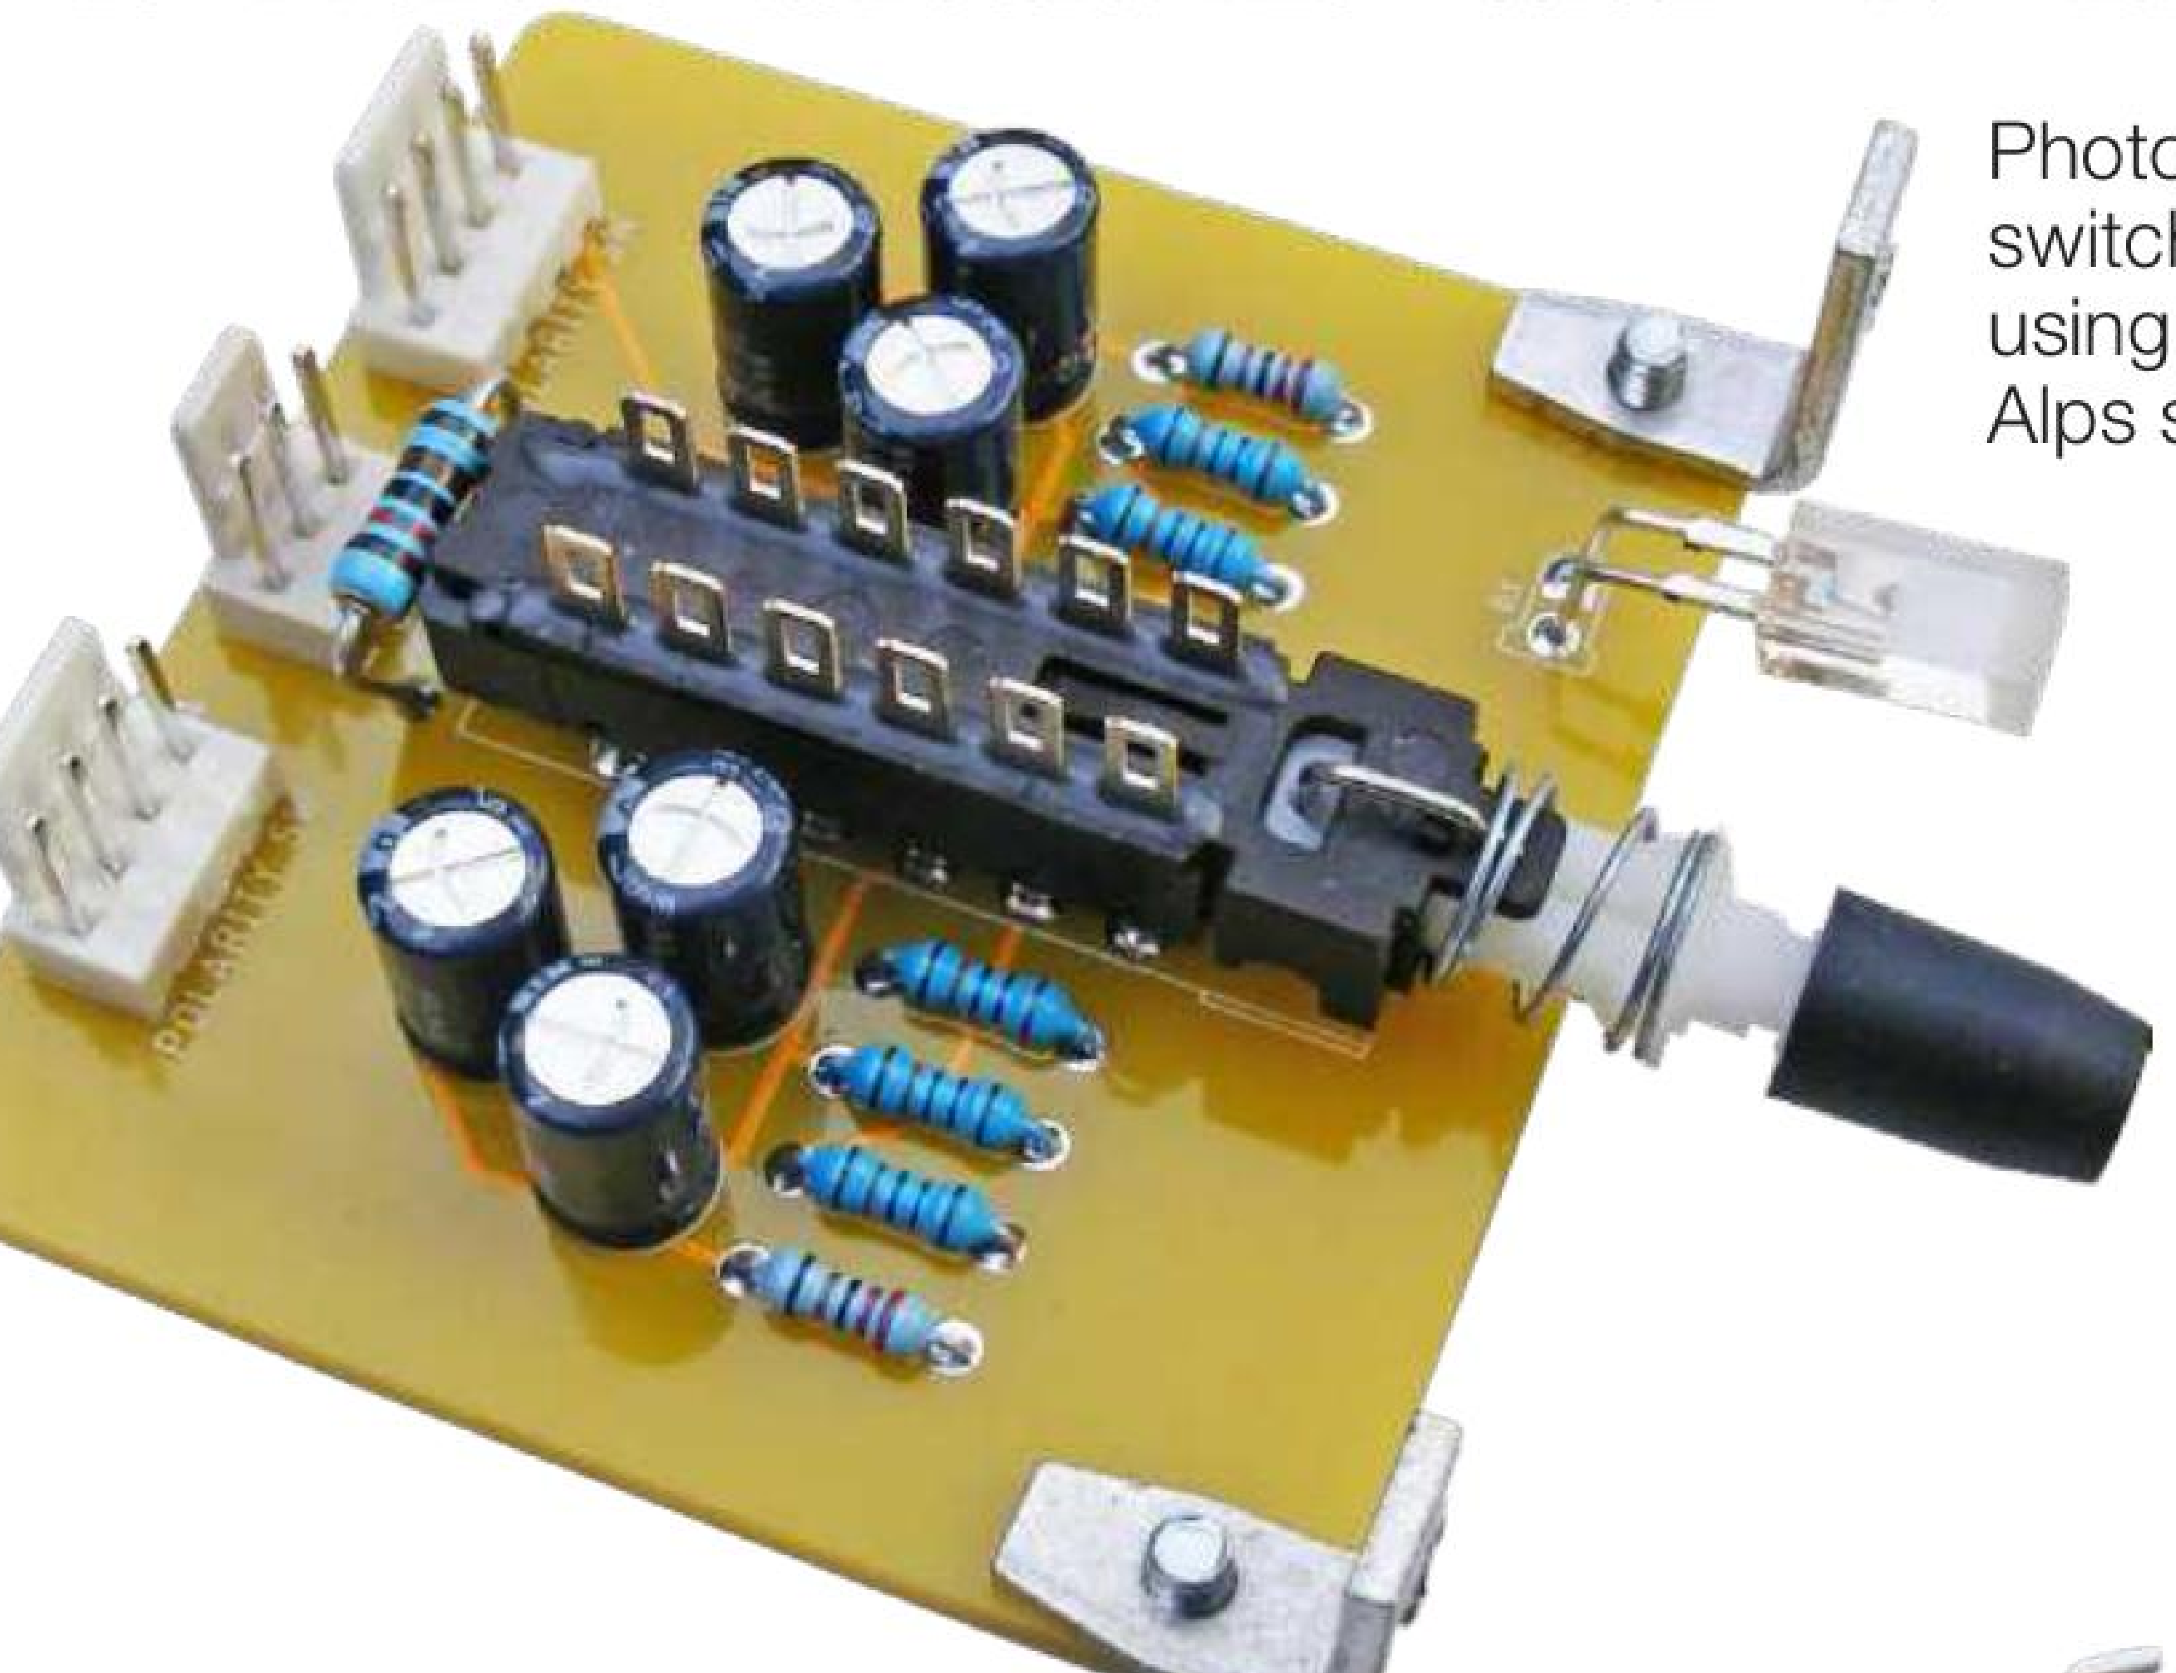

Photo 26: the polarity switch assembly using an older-style Alps switch.

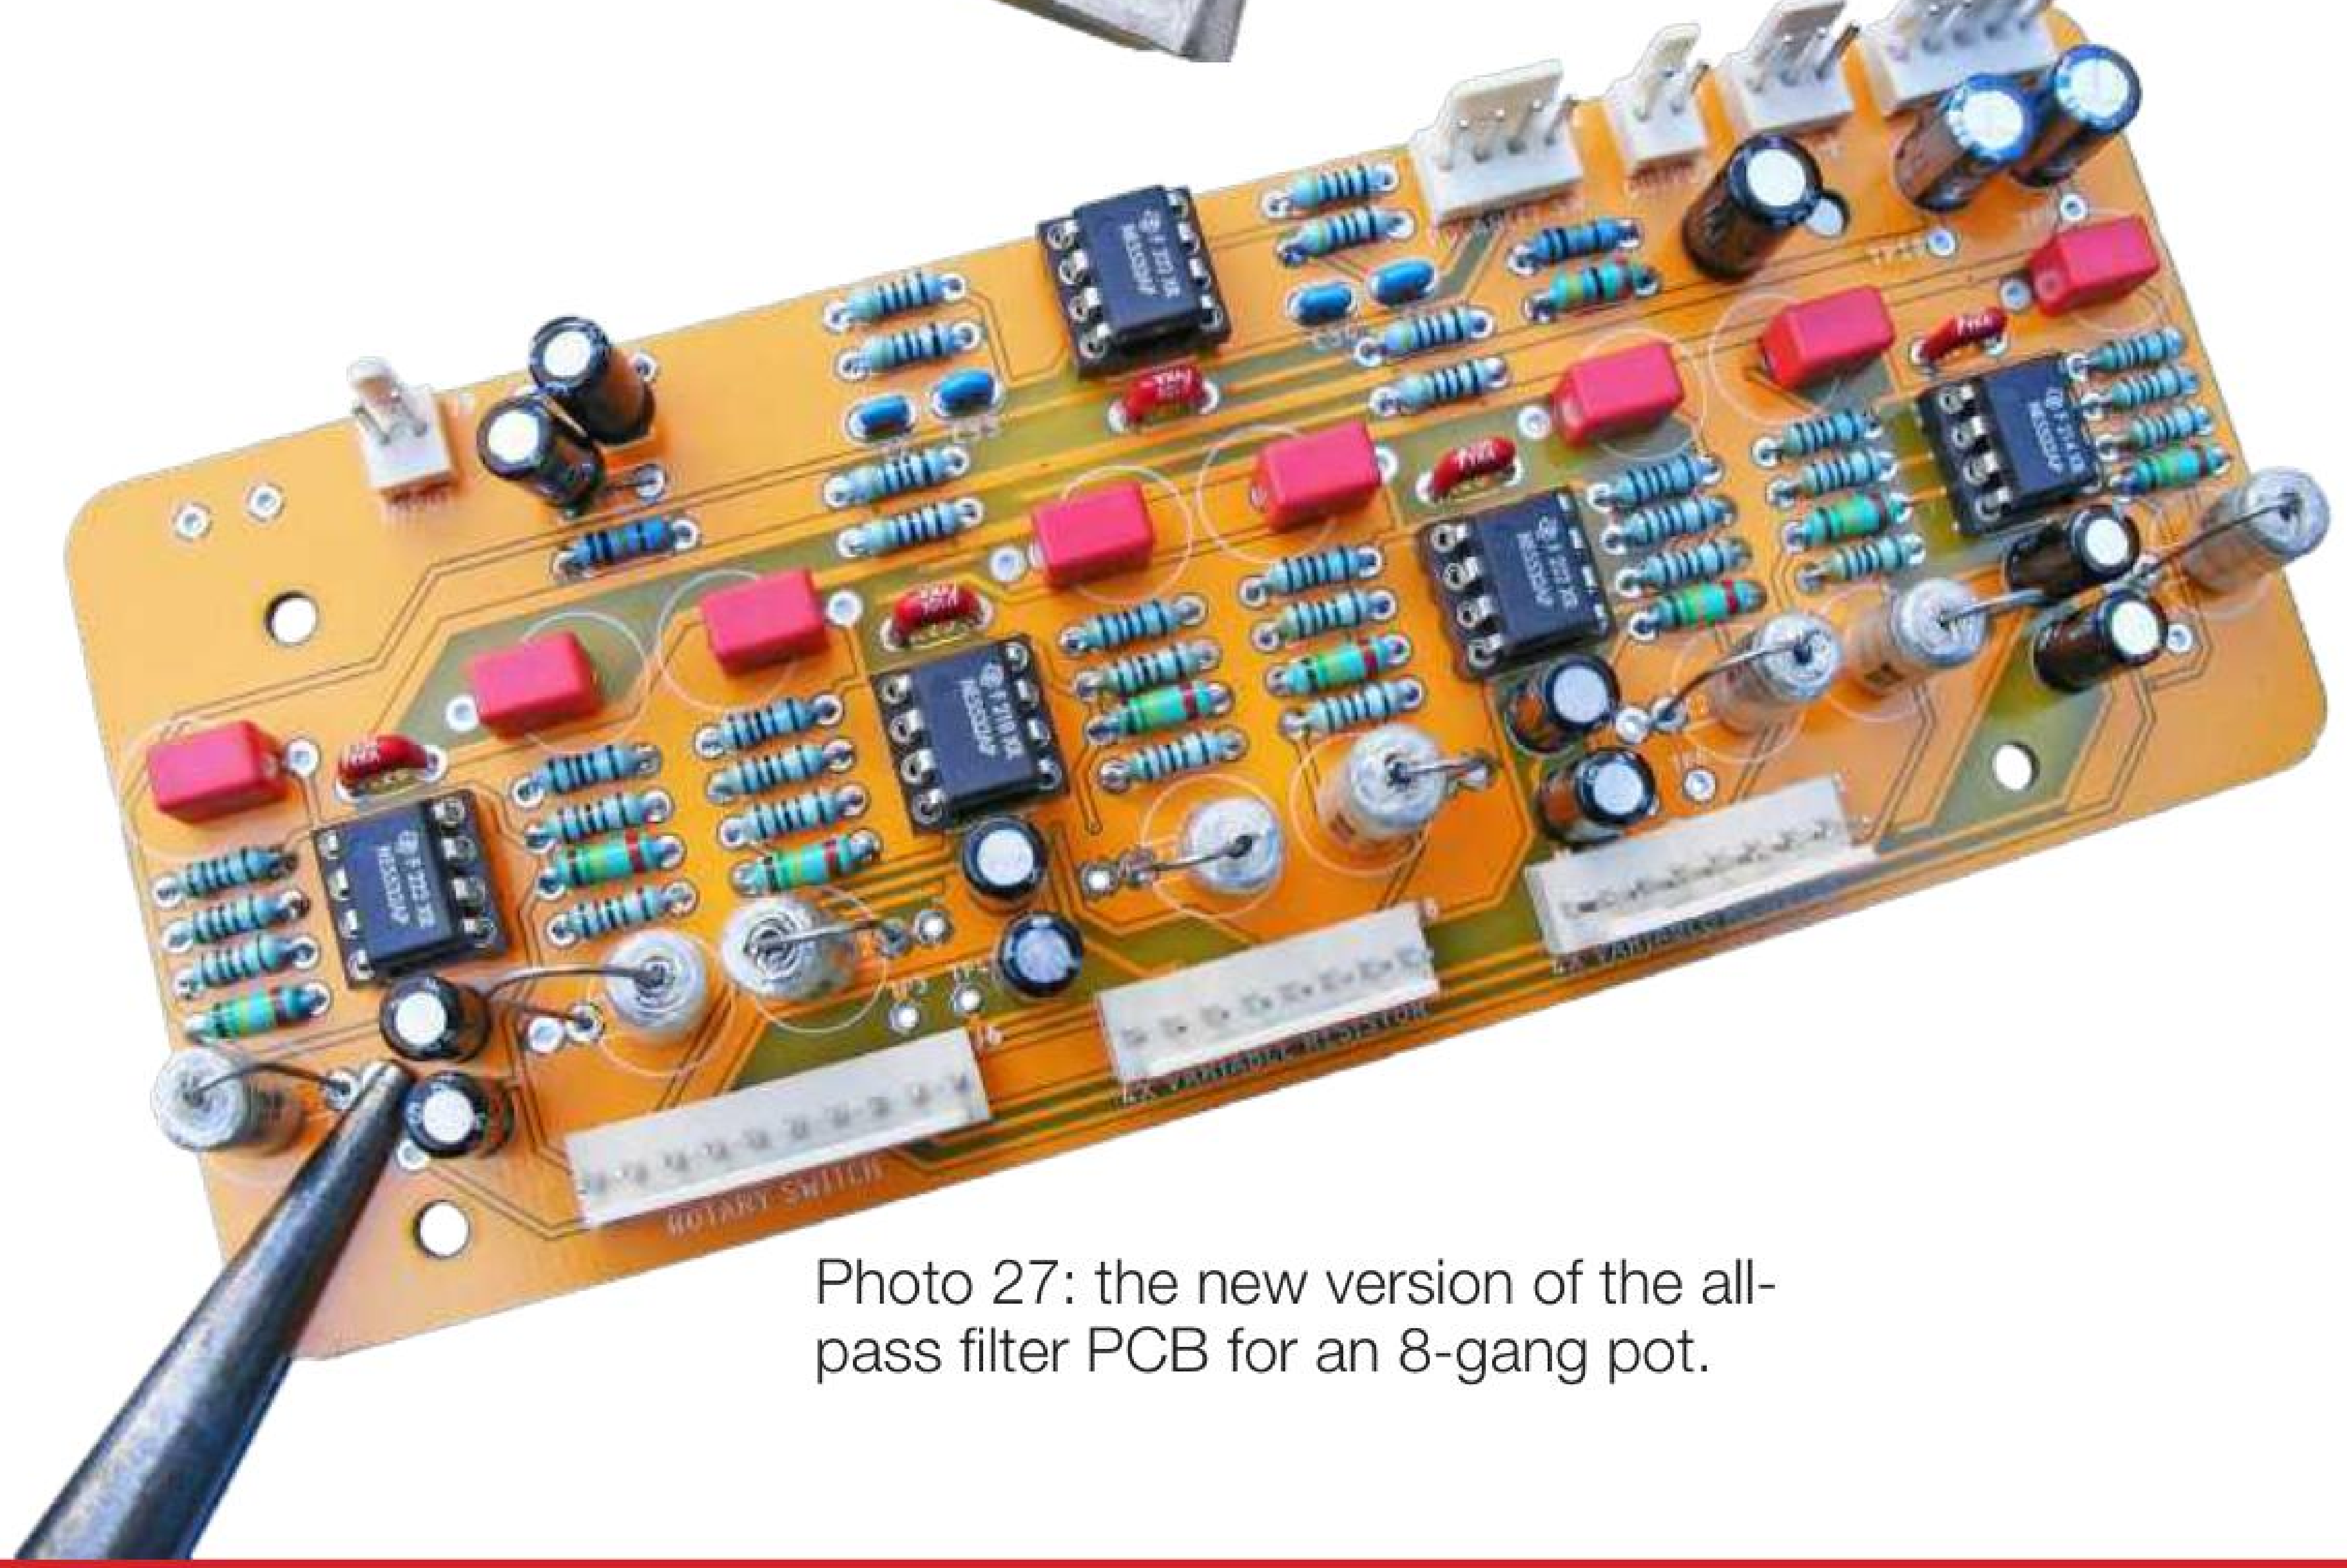

Photo 27: the new version of the all-pass filter PCB for an 8-gang pot.

Its ground current is kept separate from the signal ground to prevent clicks on switching. This connection should be taken to the 'dirty ground' on the power supply capacitors' Earth node if possible, although it is fine going to the ground sides of decoupler capacitors C30 and C31 on the main PCB.

There is also a spare changeover switch section that can be accessed from the tags on top of the switch (with the older Paxolin top Alps switches). The board also acts as a convenient way of mounting those Alps Alpine (manufacturer's part number SPUJ193700) style push switches to the front panel with a couple of fixing holes/angle brackets.

Another board is available for an alternative switch from E-switch, as shown in Photo 28. This is available for £1.60 from Mouser (PBH6UEENAG1RBLK).

These switches can be difficult to find nowadays. I used to get them from Rapid, but after ploughing through their 50 pages of industrial push-buttons, I couldn't locate them. Luckily, I still have a bag of 100 older types, so I can supply them via AOShop. Farnell and CPC still sell them; CPC's are the cheapest, while RS is mega-expensive.

It is a good idea to add a bypass switch as well. At the moment, only the input is balanced, so the circuit in Fig.37 should be used. To make it fully compatible with balanced professional broadcast gear, a balanced output with a hard bypass arrangement would be needed.

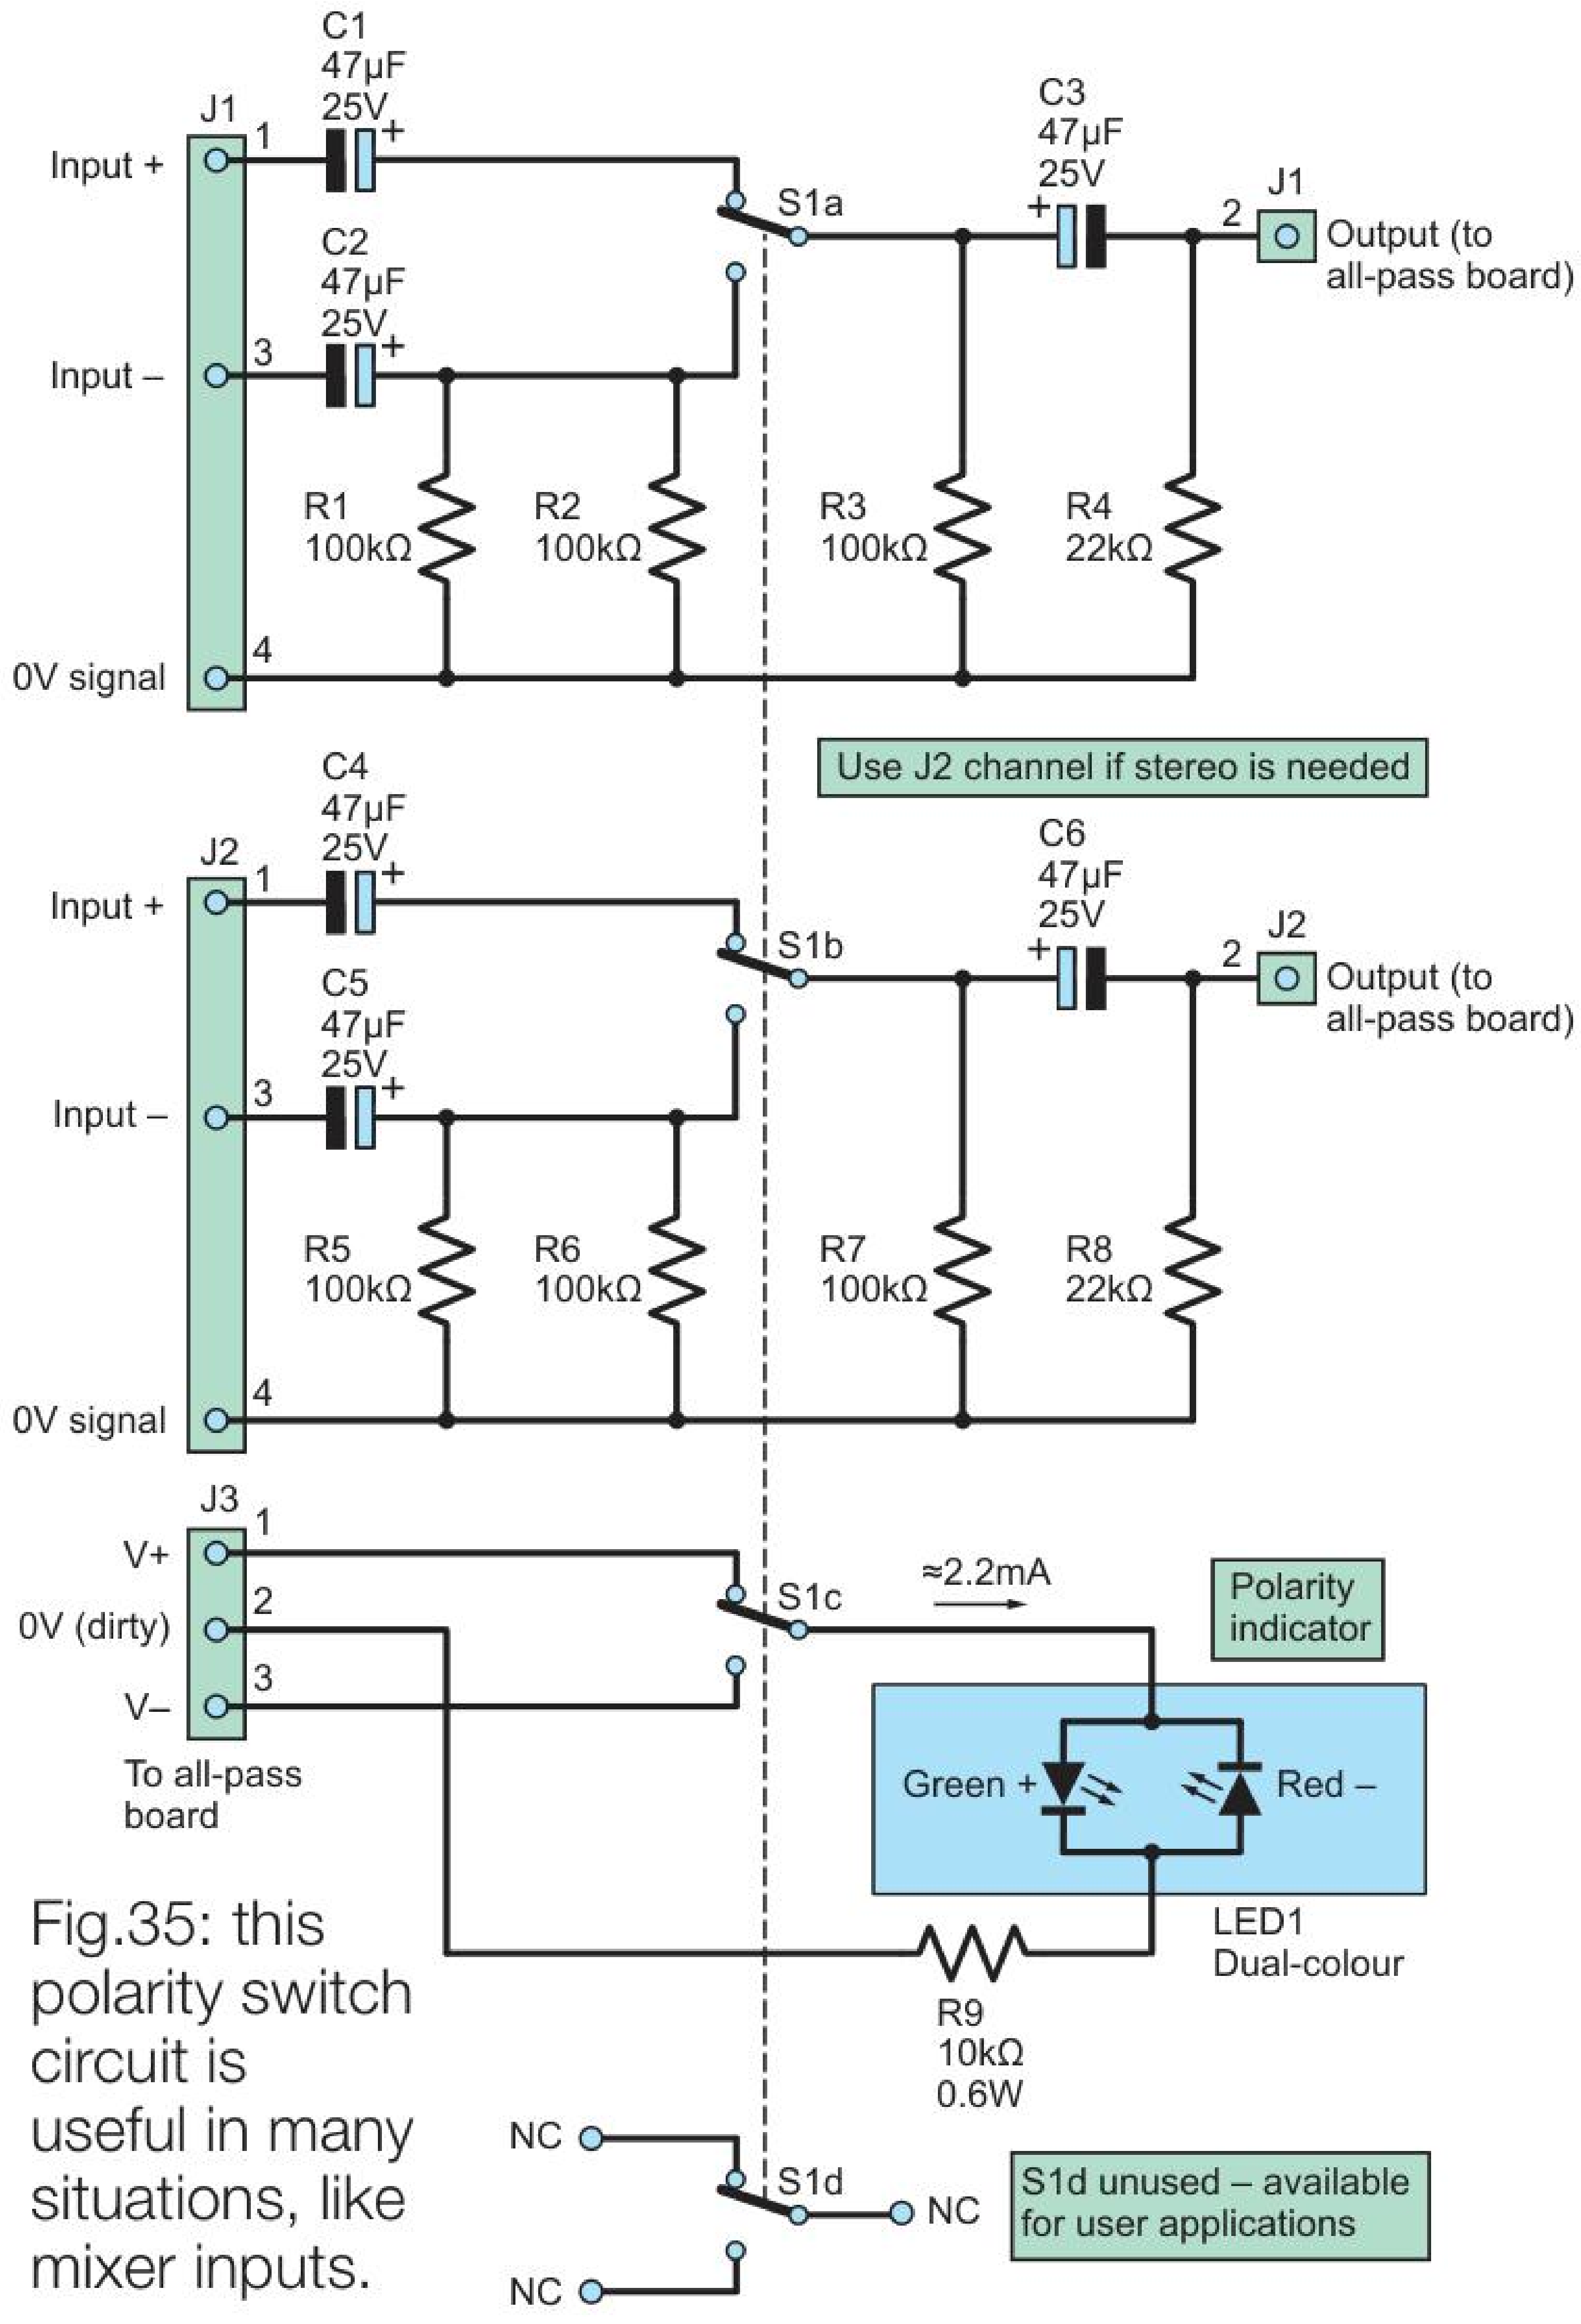
C1
47µF
25V
J1
Input +
1
C2
47µF
25V
Input –
3
R1
100kΩ
R2
100kΩ
0V signal
4
S1a
C3
47µF
25V
2
J1
Output (to all-pass board)
R3
100kΩ
R4
22kΩ
Use J2 channel if stereo is needed
C4
47µF
25V
J2
C5
47µF
25V
R5
100kΩ
R6
100kΩ
S1b
C6
47µF
25V
J2
Output (to all-pass board)
R7
100kΩ
R8
22kΩ
J3
V+
0V (dirty)
V–
To all-pass board
S1c
≈2.2mA
Polarity indicator
Green +
Red –
LED1
Dual-colour
R9
10kΩ
0.6W
NC
S1d
NC
NC
S1d unused – available for user applications

Fig.35: this polarity switch circuit is useful in many situations, like mixer inputs.

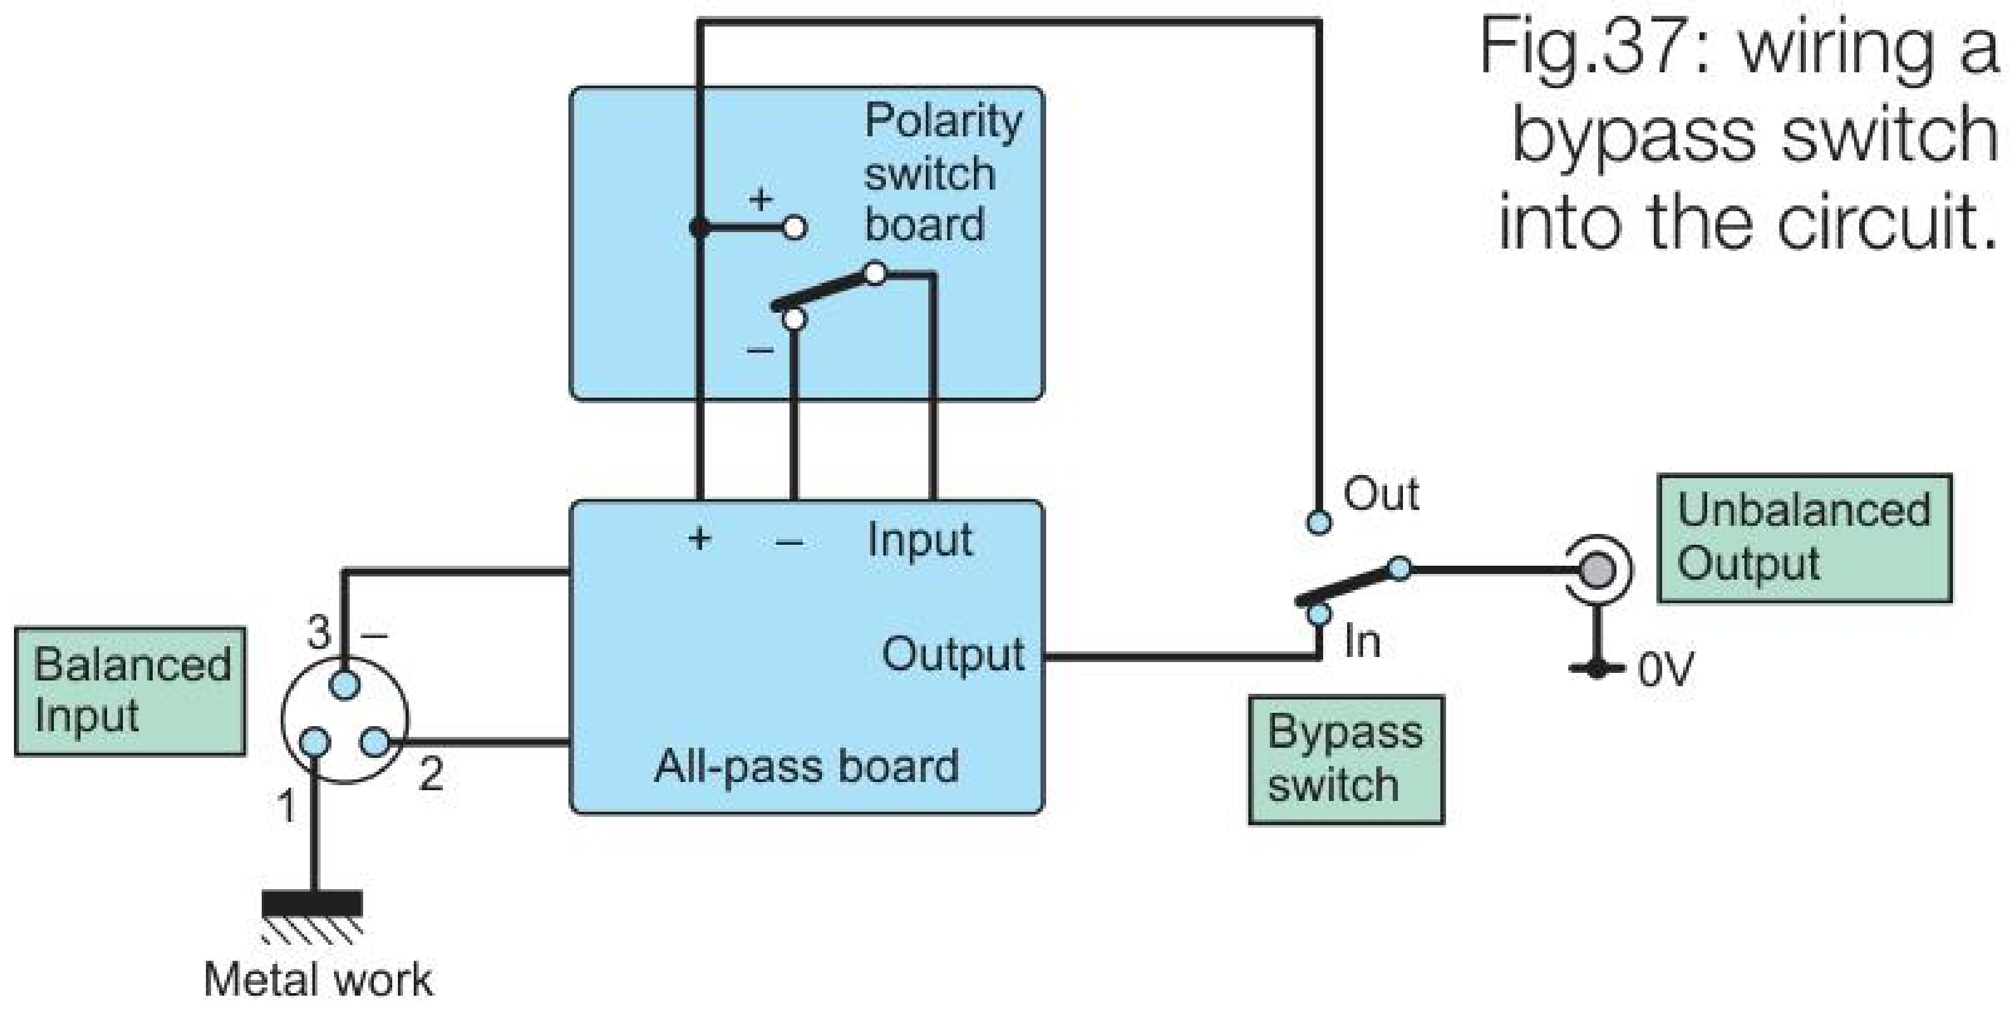
Polarity switch board
+
–
Input
Output
All-pass board
Balanced Input
3
1
2
Metal work
Out
In
Bypass switch
Unbalanced Output
0V

Fig.37: wiring a bypass switch into the circuit.

All-pass effects

All-pass filtering can be an effect in itself, giving an illusion of extra bass, a characteristic 'chesty' booming sound on male speech. This effect is also produced by small bass-reflex speakers with 6th-order bass roll-offs, typified by most small active studio monitors such as those by Genelec, Adams and KRK.

From microphone, to preamp, to equalisation, to AC-coupled audio interface, to the amplifier, the sum of all the phase shifts due to the culminating high-pass poles can be quite destructive to bass coherence.

Electronically, a total of around 16 all-pass poles with 100nF capacitors are needed for this psychoacoustic delay-based 'bass boost' to be audible, but this is testimony to how insensitive to phase human hearing is. Compare this to our perception of frequency, where just one or two poles of filtering are clearly audible.

Phase shift can be modulated by a low-frequency oscillator or LFO, giving an effect similar to vibrato, but it has a strangely nausea-inducing quality. In stereo, it can almost give rise to seasickness! That's because it isn't pitch but frequency modulation, which demolishes harmonic relationships.

Dedicated frequency shifters produce a static version of this effect. Harvey (my son) and I have been messing about with all-pass filtering in music software. Although a research rabbit hole, this can lead to insights that can then be made into analog products.

New eight-gang pot version

Since the Editor pointed out that eight-gang pots were available from Bourns via Mouser, it seemed desirable to have the all-pass board redesigned for them, along with a new pot mounting board. It's a bit silly that I had never found them, since I've just noticed they were used in John Clarke's active crossover, back in September 2018, and I've been needing some for years.

The part number is PTD908-2015K-C503, where the C503 means 50kΩ reverse/anti log. The new all-pass circuit is the same as before, except that four extra 47µF pot 'anti-scratch' coupling capacitors, C35 to C38, have been added.

Originally these pot coupling capacitors (C6, C8, C11 and C14) were 10µF, but I found some 47µF 16V bi-polar types that fitted fine. It's always worth going bigger and bipolar if you can. There are also an additional four 150kΩ DC path resistors, R41 to R44. Basically, it's

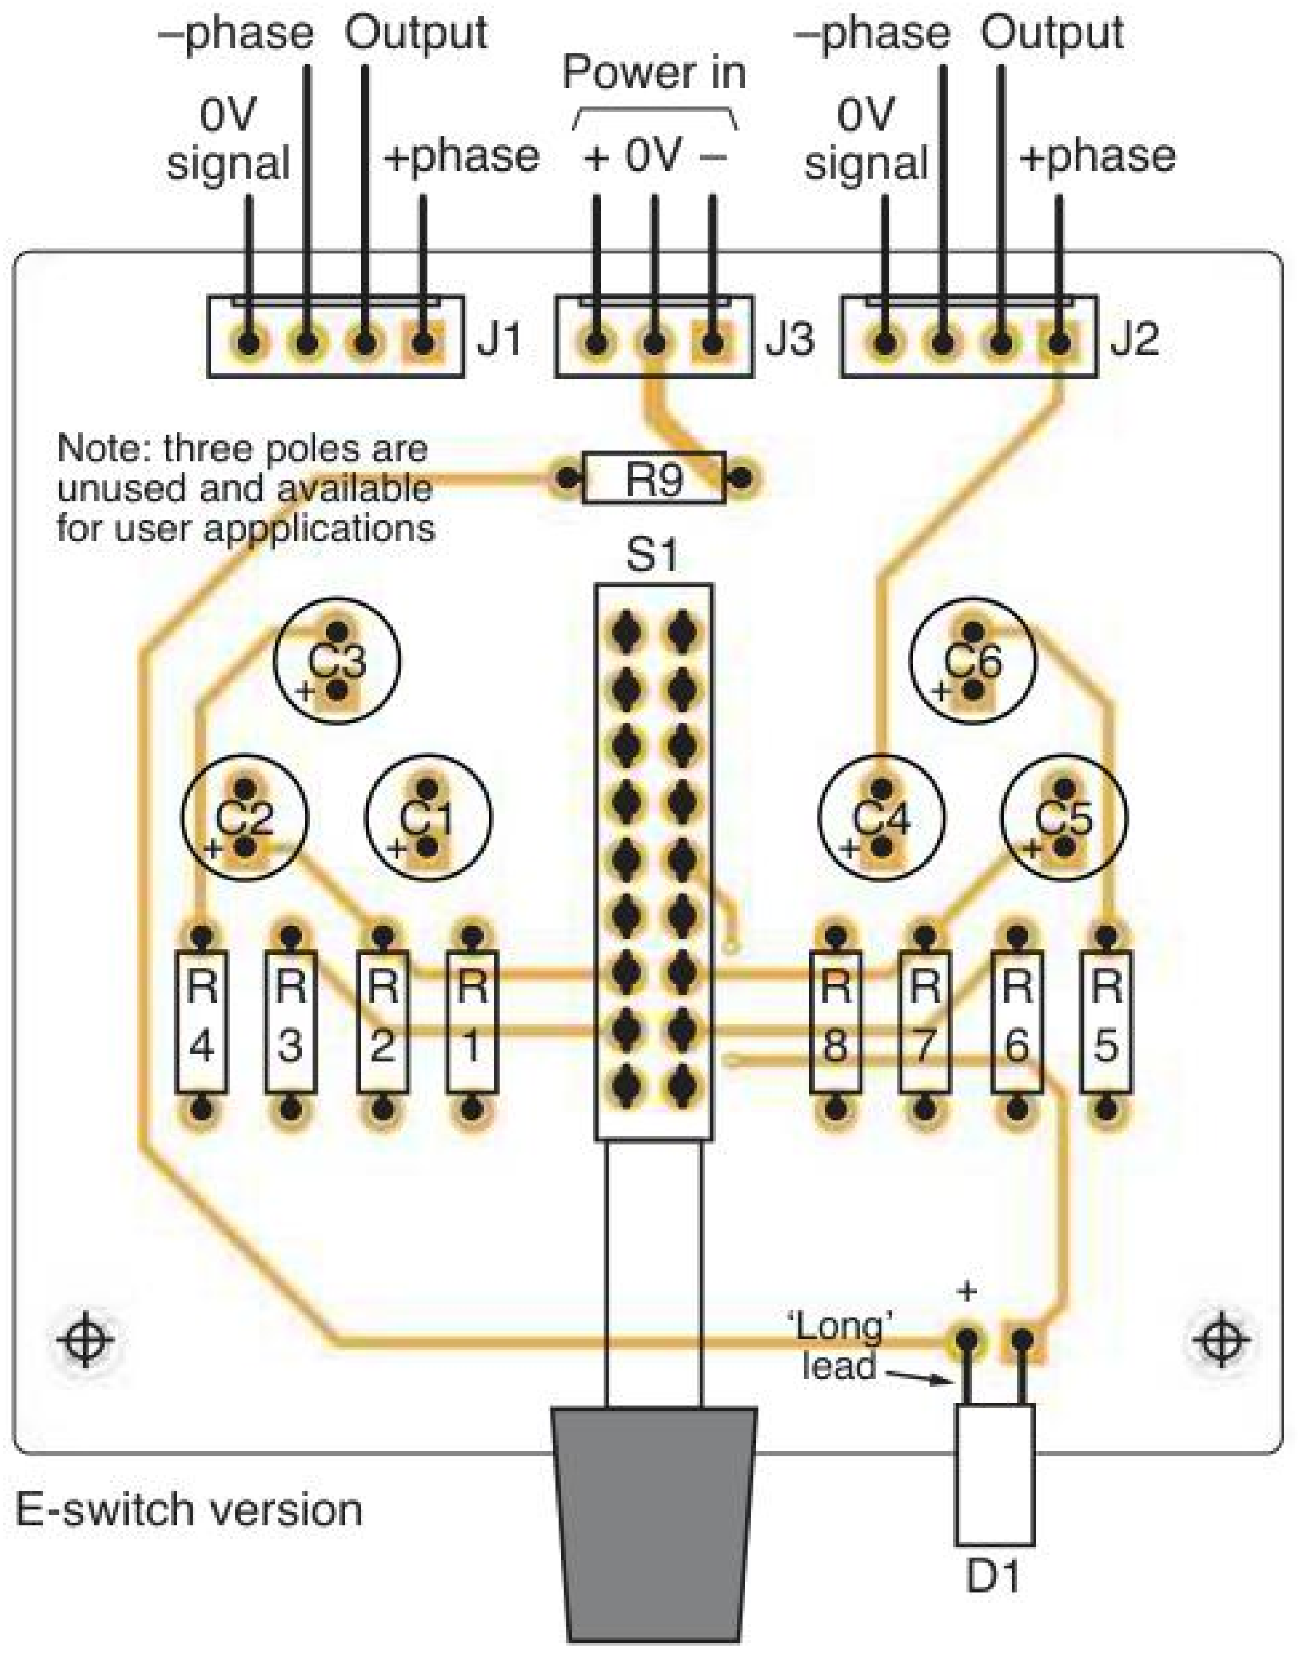
–phase
Output
0V signal
+phase
Power in
+ 0V –
–phase
Output
0V signal
+phase
J1
J3
J2
Note: three poles are unused and available for user appplications
R9
S1
C3
C6
C2
C1
C4
C5
R4
R3
R2
R1
R8
R7
R6
R5
'Long' lead
+
D1

E-switch version

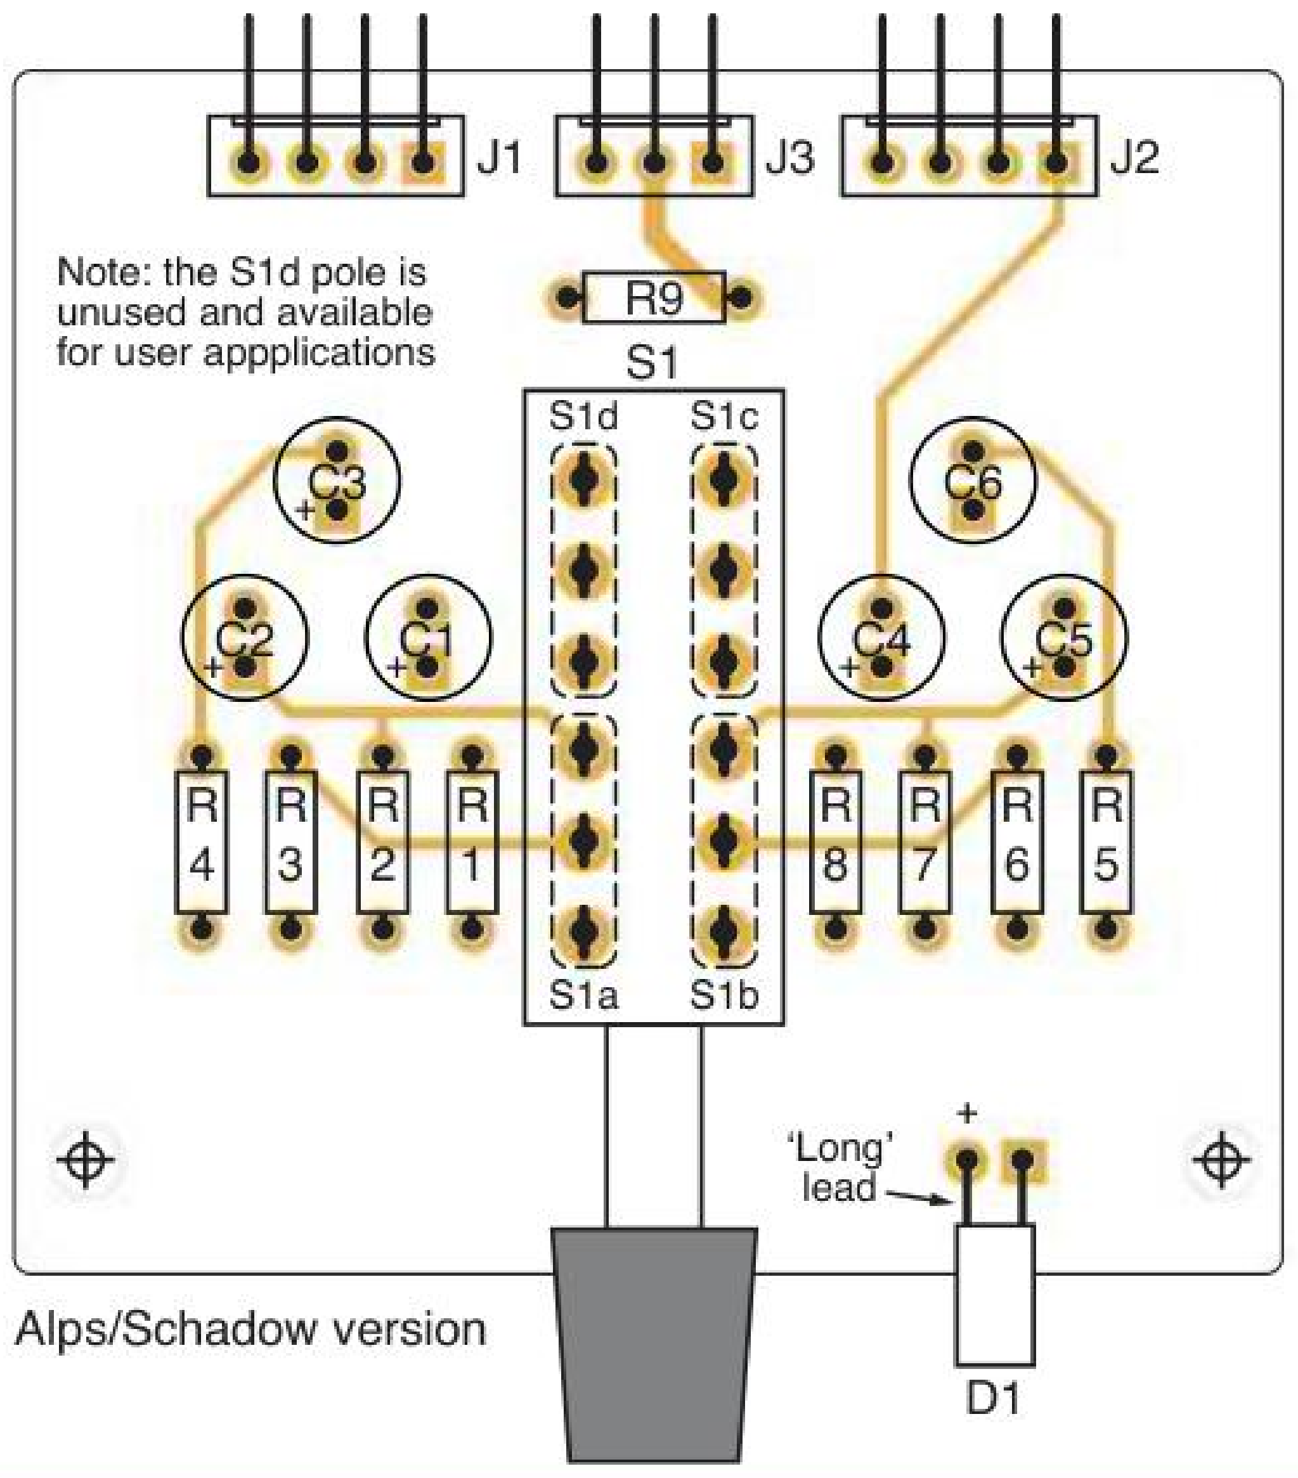
J1
J3
J2
Note: the S1d pole is unused and available for user appplications
R9
S1
S1d
S1c
S1a
S1b
C3
C6
C2
C1
C4
C5
R4
R3
R2
R1
R8
R7
R6
R5
'Long' lead
+
D1

Alps/Schadow version

◀ Fig.36: the polarity switch PCBs. The boards plug into the eight-gang pot board.

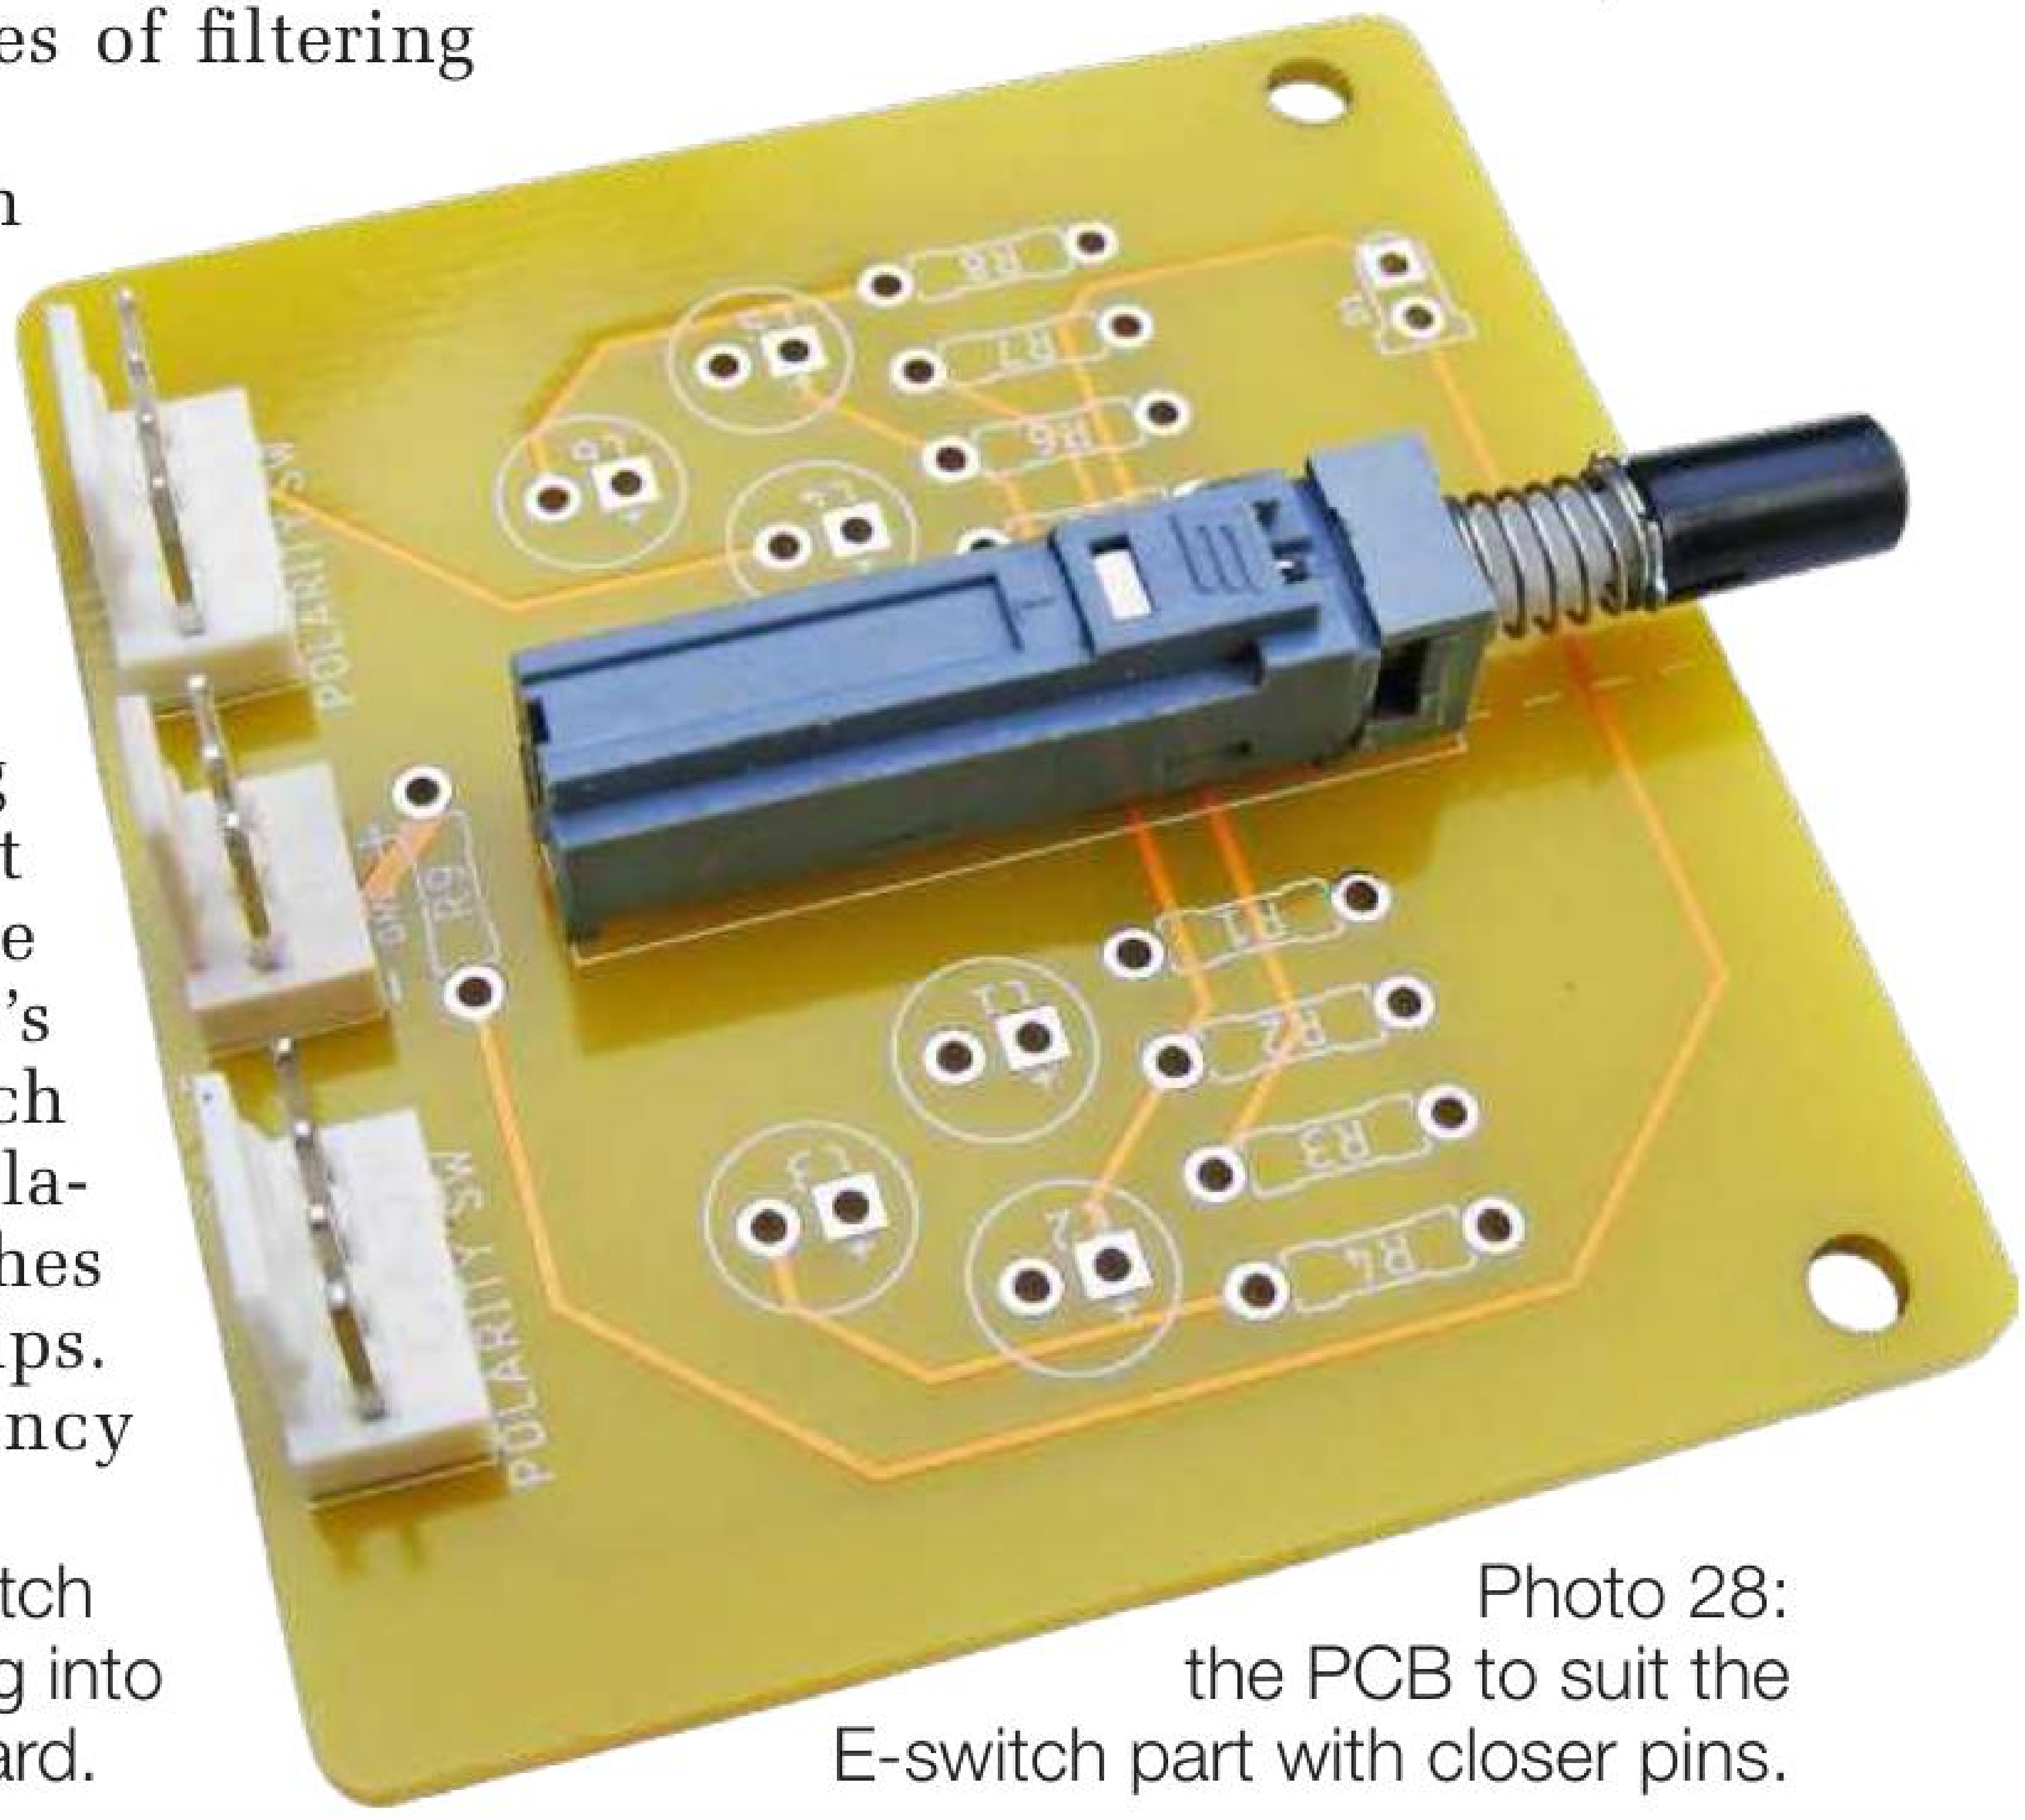
POLARITY SW
POLARITY SW
R9
R1
R2
R3
R6
R7

Photo 28: the PCB to suit the E-switch part with closer pins.

a duplication of the original circuitry around IC2 and IC3, and making it the same around IC4 and IC5.

The new overlays are shown in Fig.38 and Fig.39, while the completed assembly can be seen in Photo 29. This board has been made with vertically mounted polystyrene capacitors, shown in detail in Photo 30.

The eight-gang pot board is also useful for stereo units and gives a wider adjustment range on a single all-pass filter board compared to the four-gang board. The assembly is shown in Photo 31. While redesigning this board, we also took the opportunity to add the missing pair of 47pF capacitors, now numbered C33 and C34.

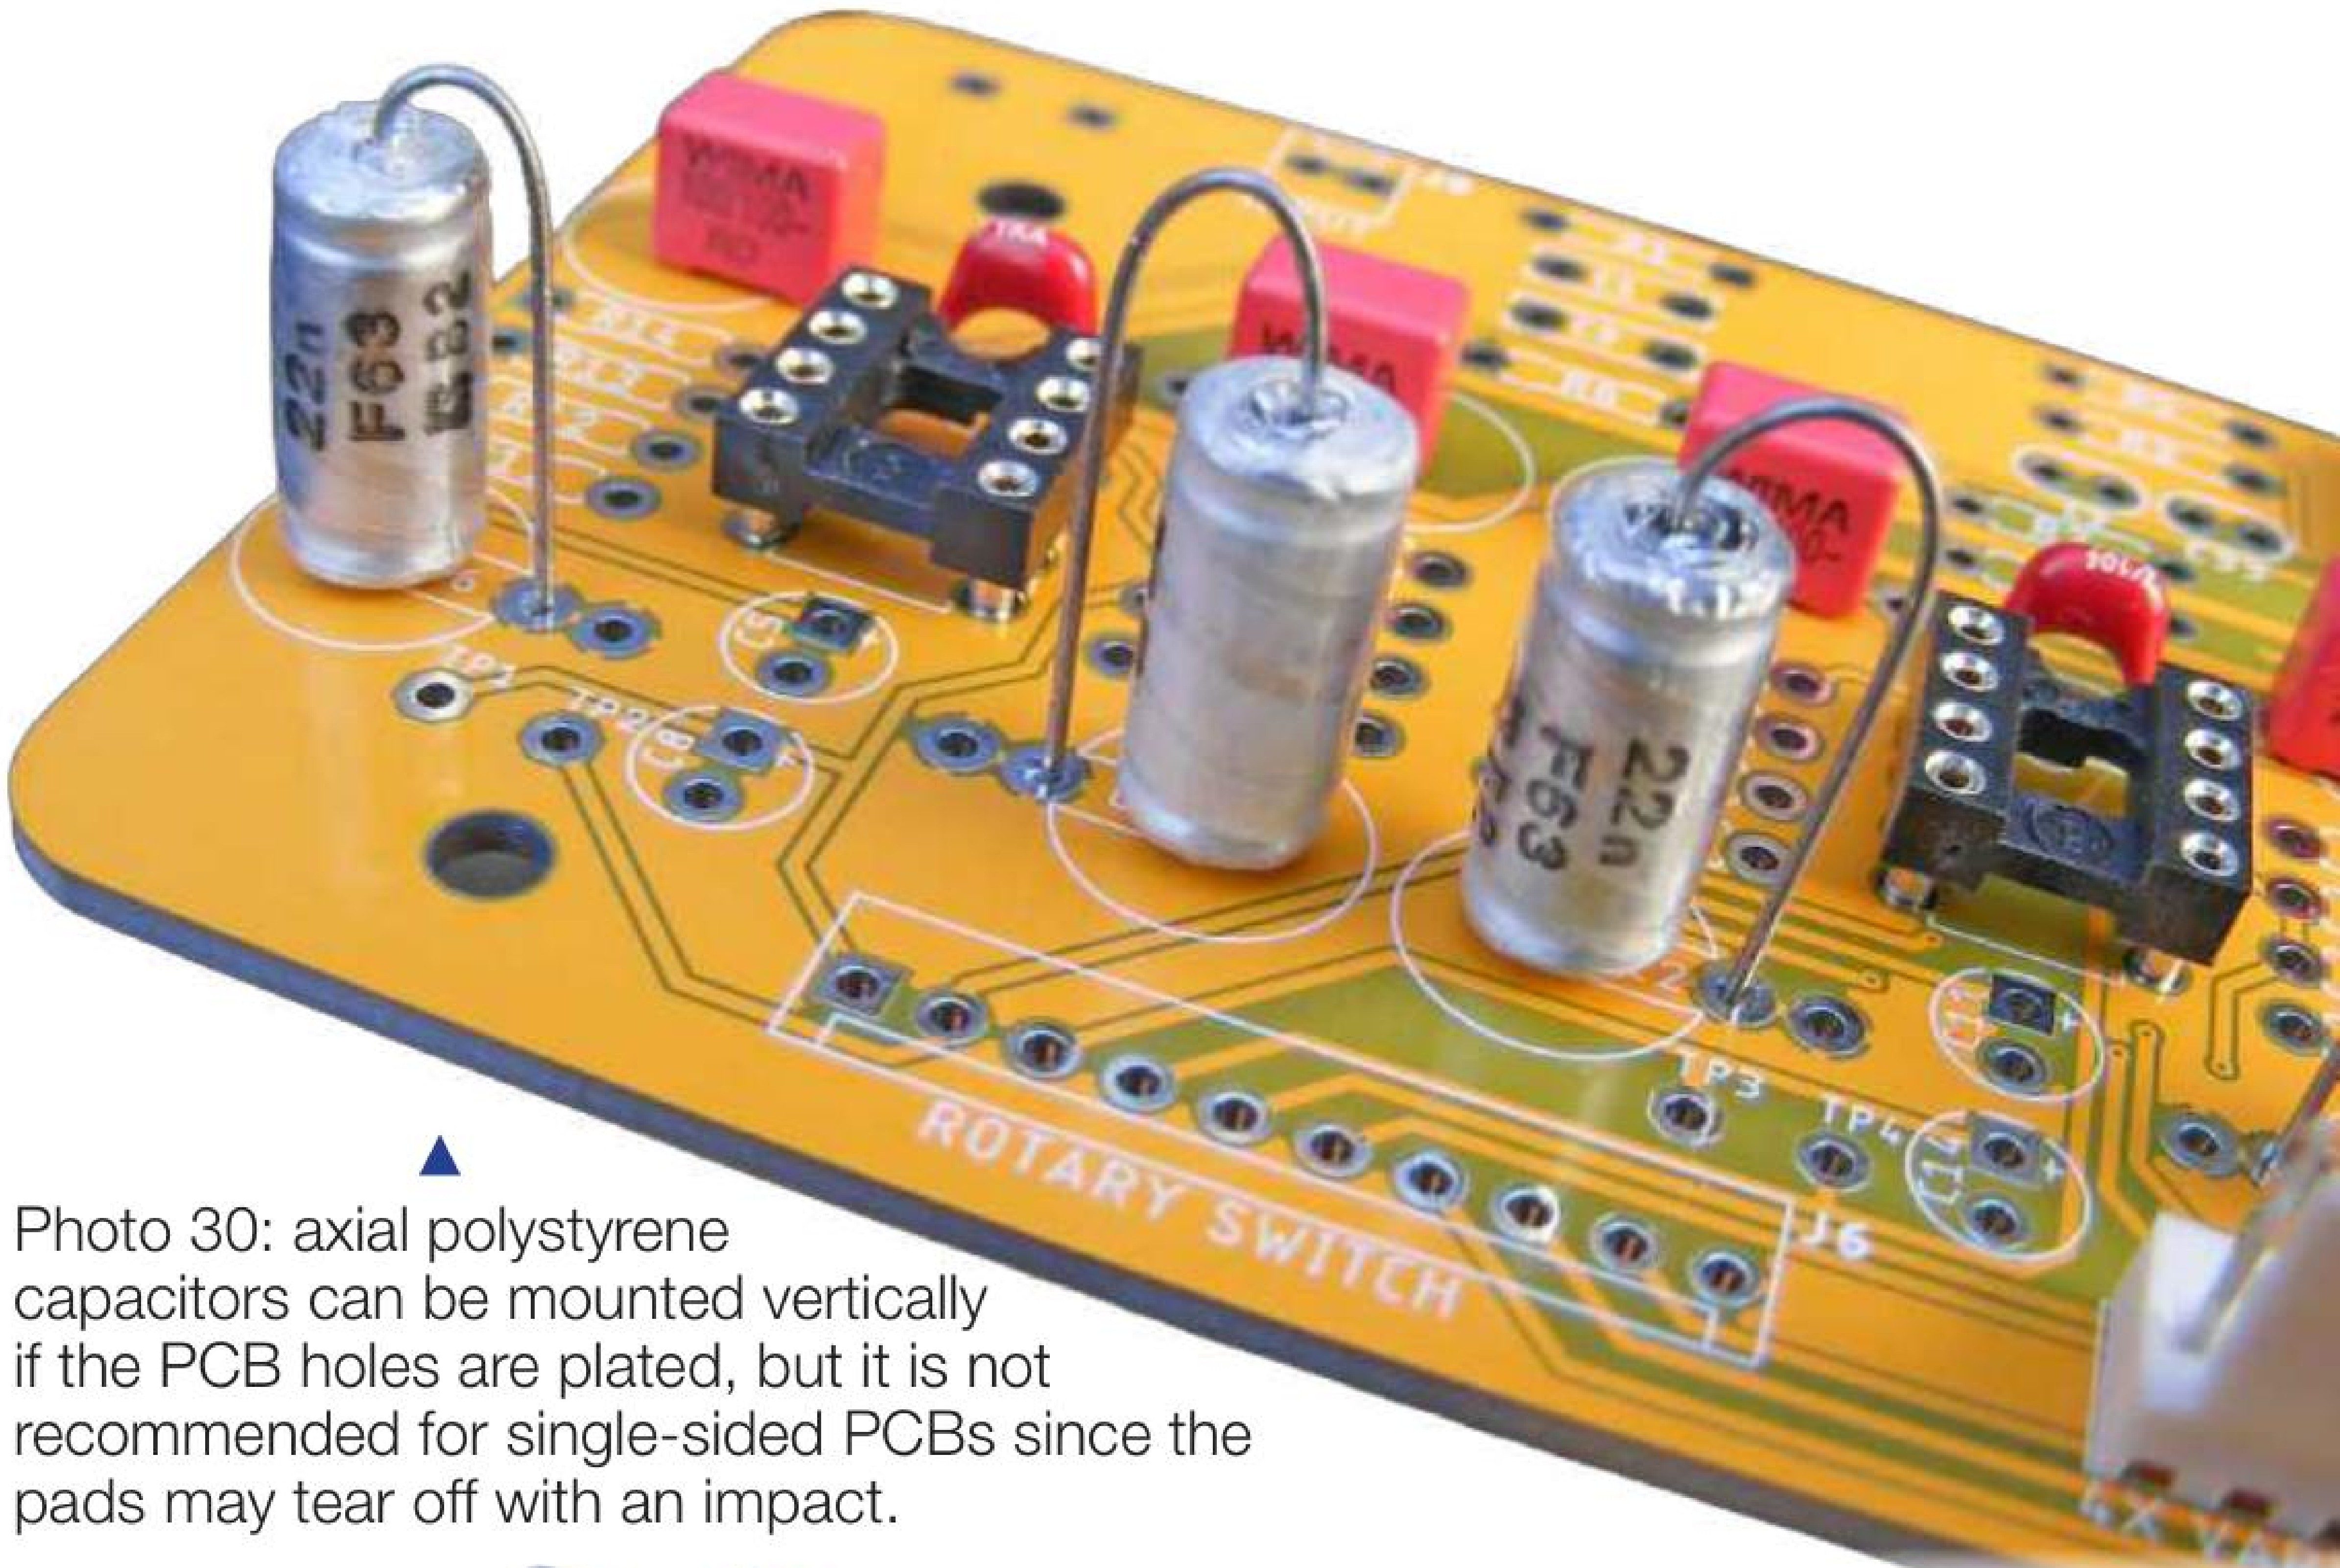

▲
Photo 30: axial polystyrene capacitors can be mounted vertically if the PCB holes are plated, but it is not recommended for single-sided PCBs since the pads may tear off with an impact.

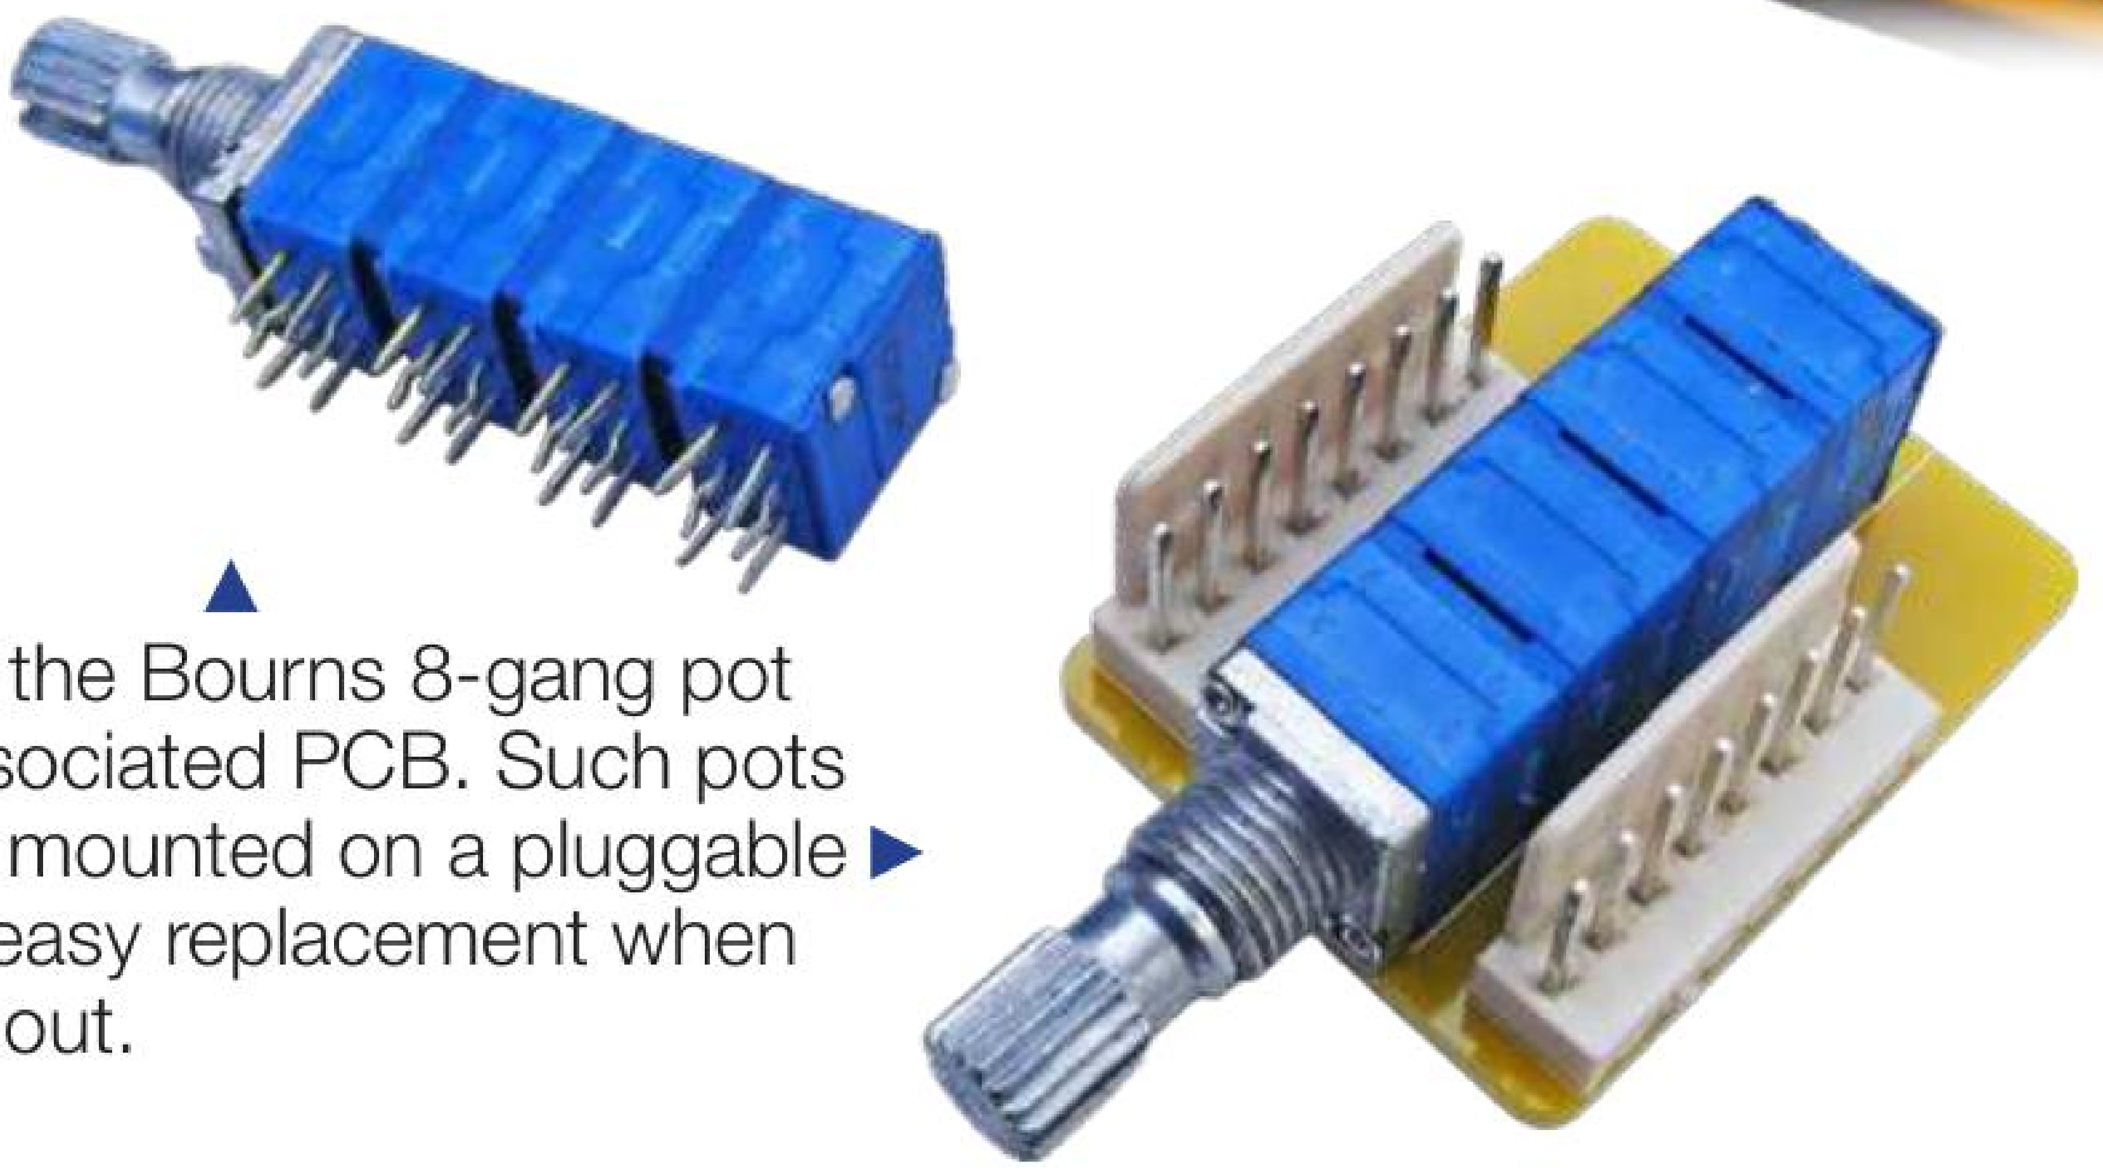

▲
Photo 31: the Bourns 8-gang pot and its associated PCB. Such pots should be mounted on a pluggable ▶ board for easy replacement when they wear out.

Testing

Start by testing it with a sinewave signal from an oscillator or waveform generator. The unit should have no effect unless you rotate the pot rapidly, where there will be a slight frequency modulation due to the varying delay.

A pulse wave or uneven mark-space ratio square wave is a good way to test that each of the delay stages is working. As the switch is incremented, the number of 'squiggles' in the waveform should increase. Photo 32 shows the effect of one all-pass stage on a square wave.

Speech waveform symmetriciser

This is a popular project among radio amateurs who need to maximise their speech intelligibility in poor reception conditions. This is achieved by scrambling the phase of the speech harmonics relative to the fundamental to make the waveform more symmetrical using an all-pass filter.

This reduces the peak-to-peak voltage, enabling higher volume, maximising the transmitter modulation. This circuit is sometimes called a phase shuffler, an ambiguous term first coined by Blumlein for his binaural stereo-widening circuit. Michael Gerzon also used the term shuffler to describe a side-signal +8dB low-frequency booster for his stereo microphone array.

The original way to process speech was by crude symmetrical diode clipping, but this causes gross distortion like a fuzzed guitar. This is minimised by telephone-style band-pass limiting using 300Hz high-pass and 3.4kHz low-pass filtering before and after the clipping. A more advanced system is to convert the audio to RF, clip it, then convert it back.

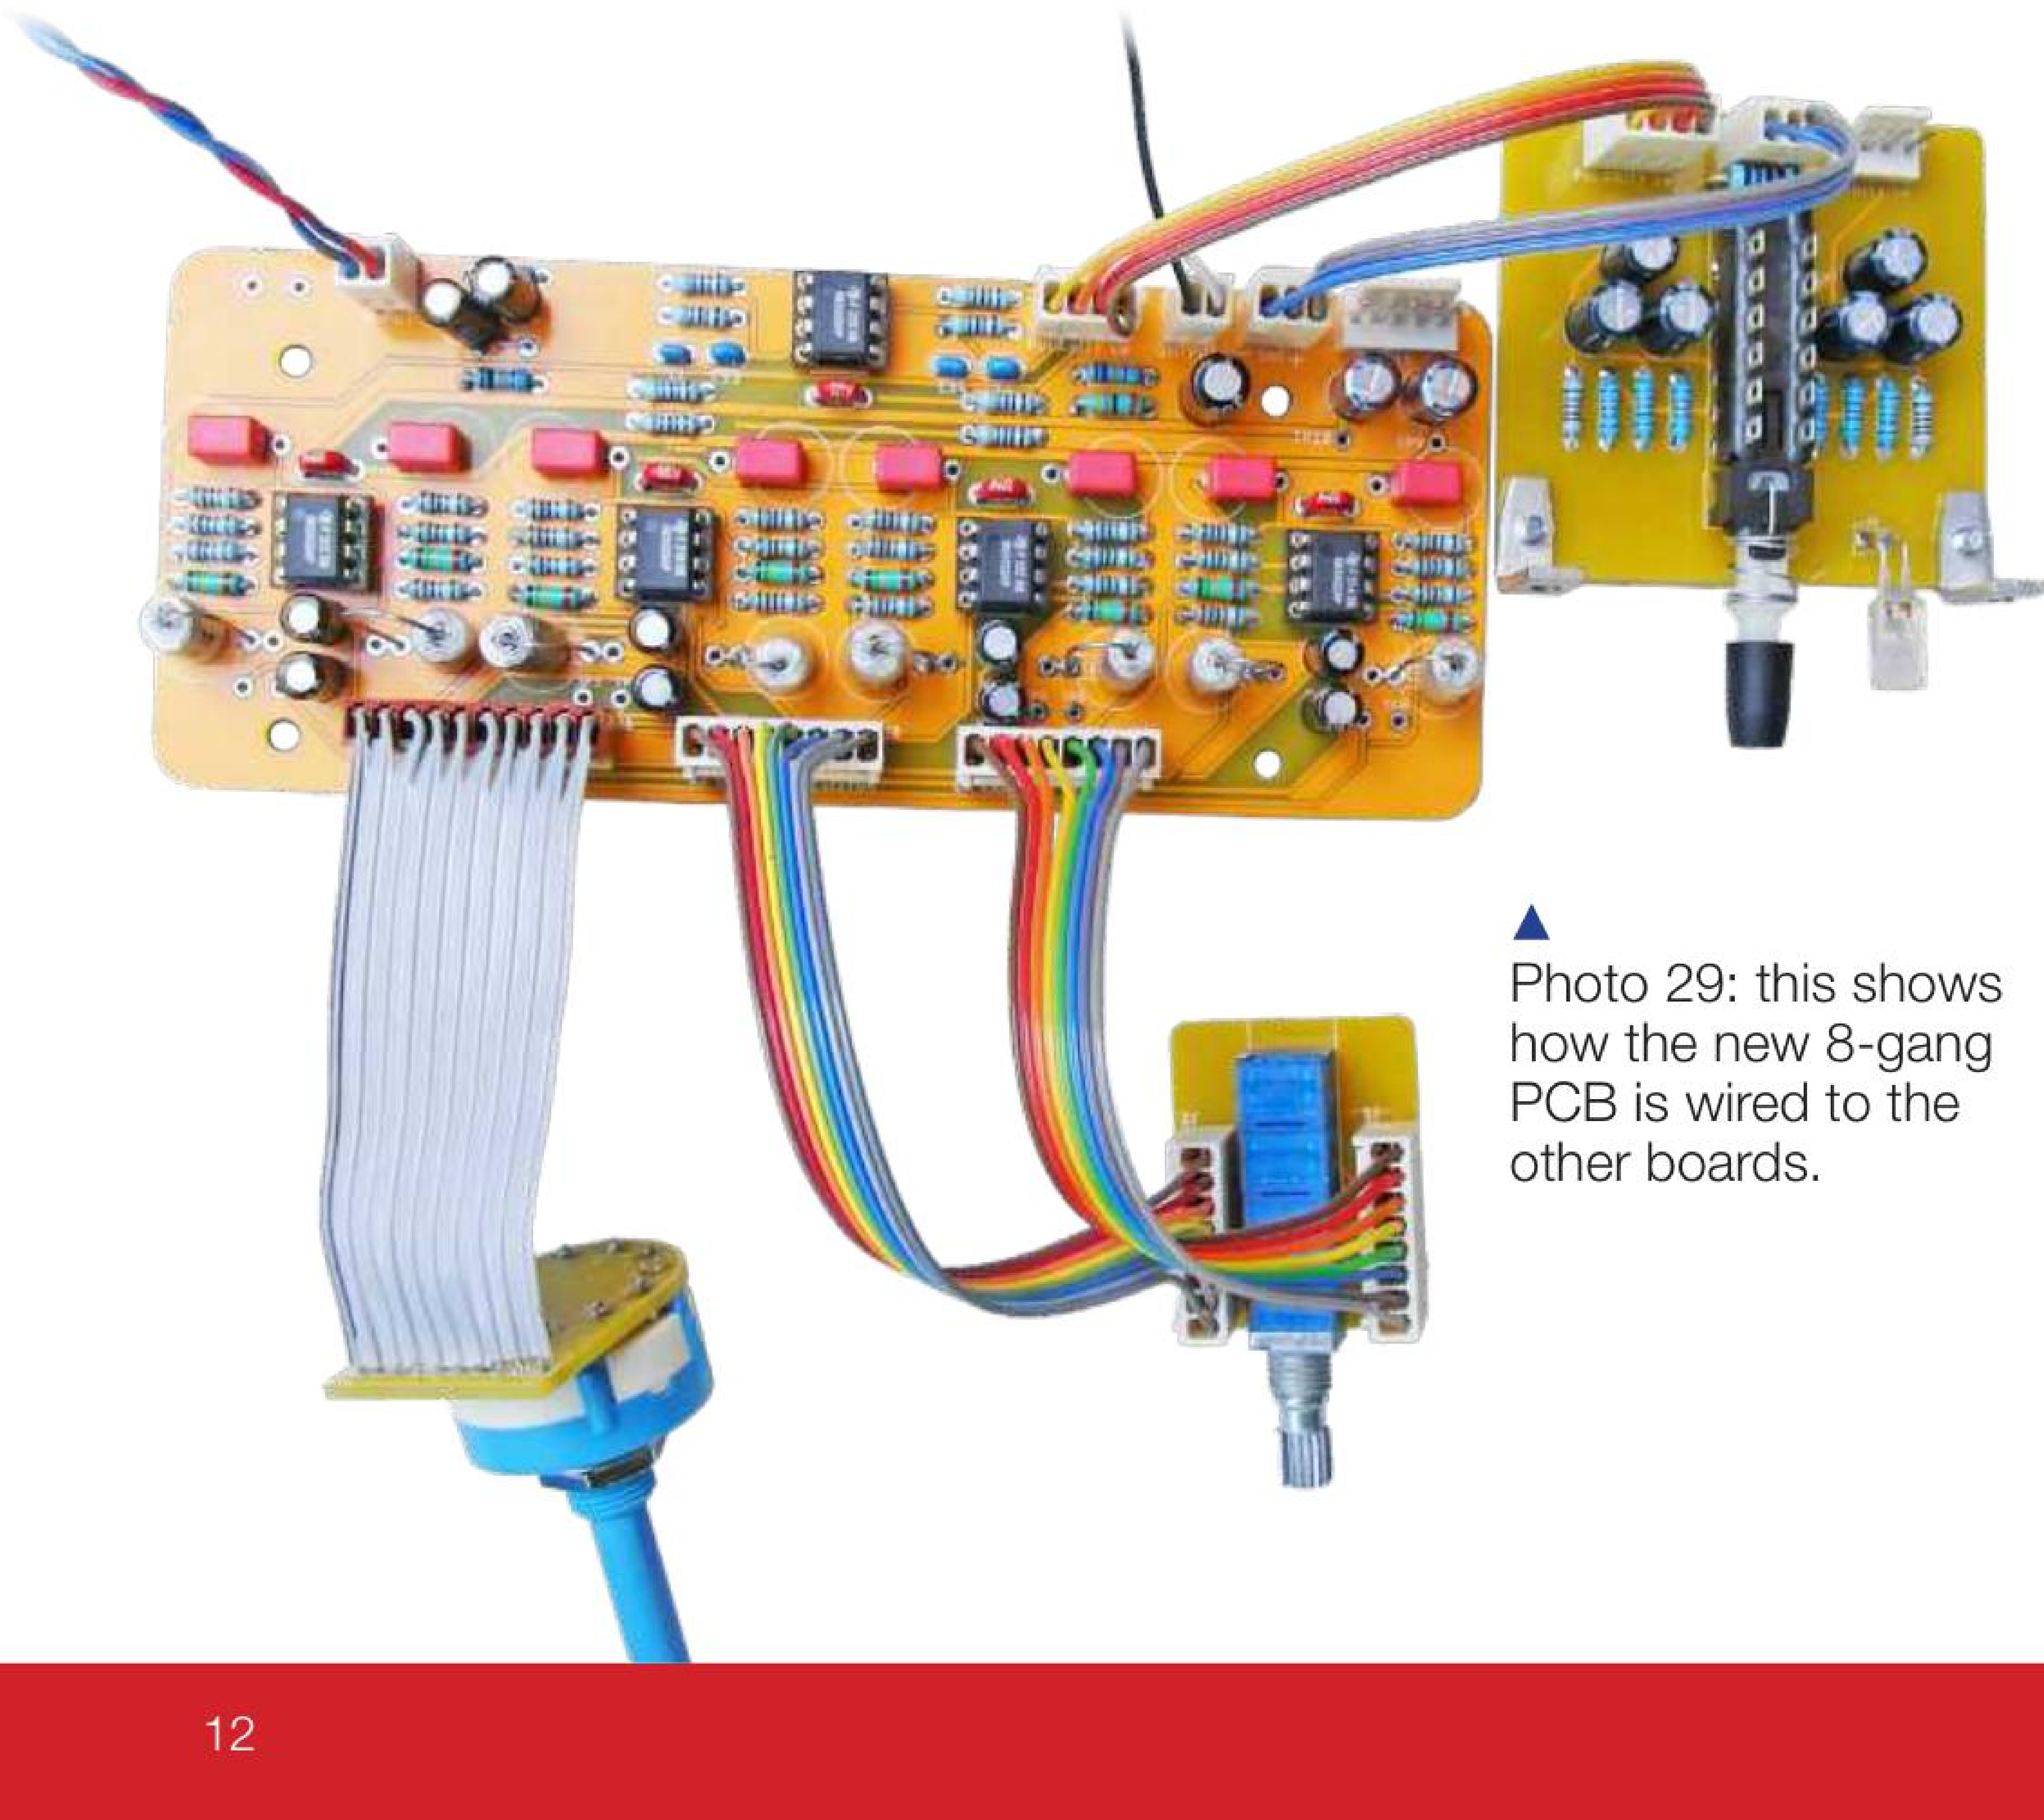

▲
Photo 29: this shows how the new 8-gang PCB is wired to the other boards.

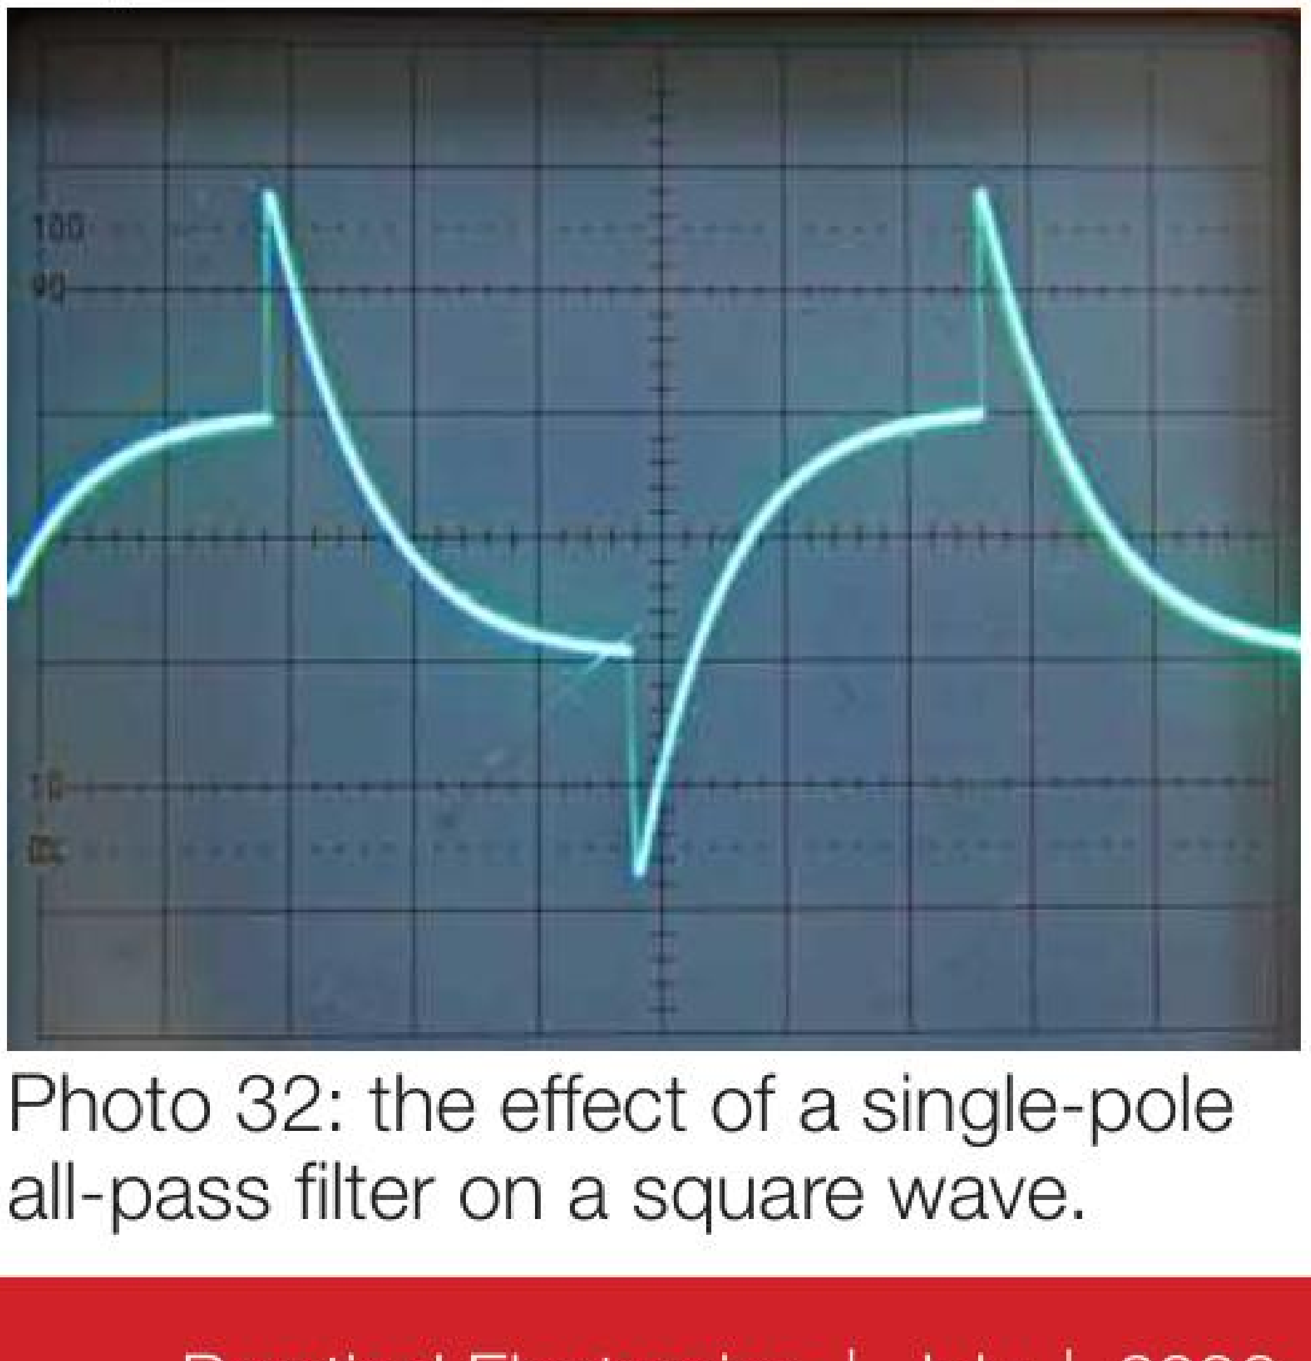

Photo 32: the effect of a single-pole all-pass filter on a square wave.

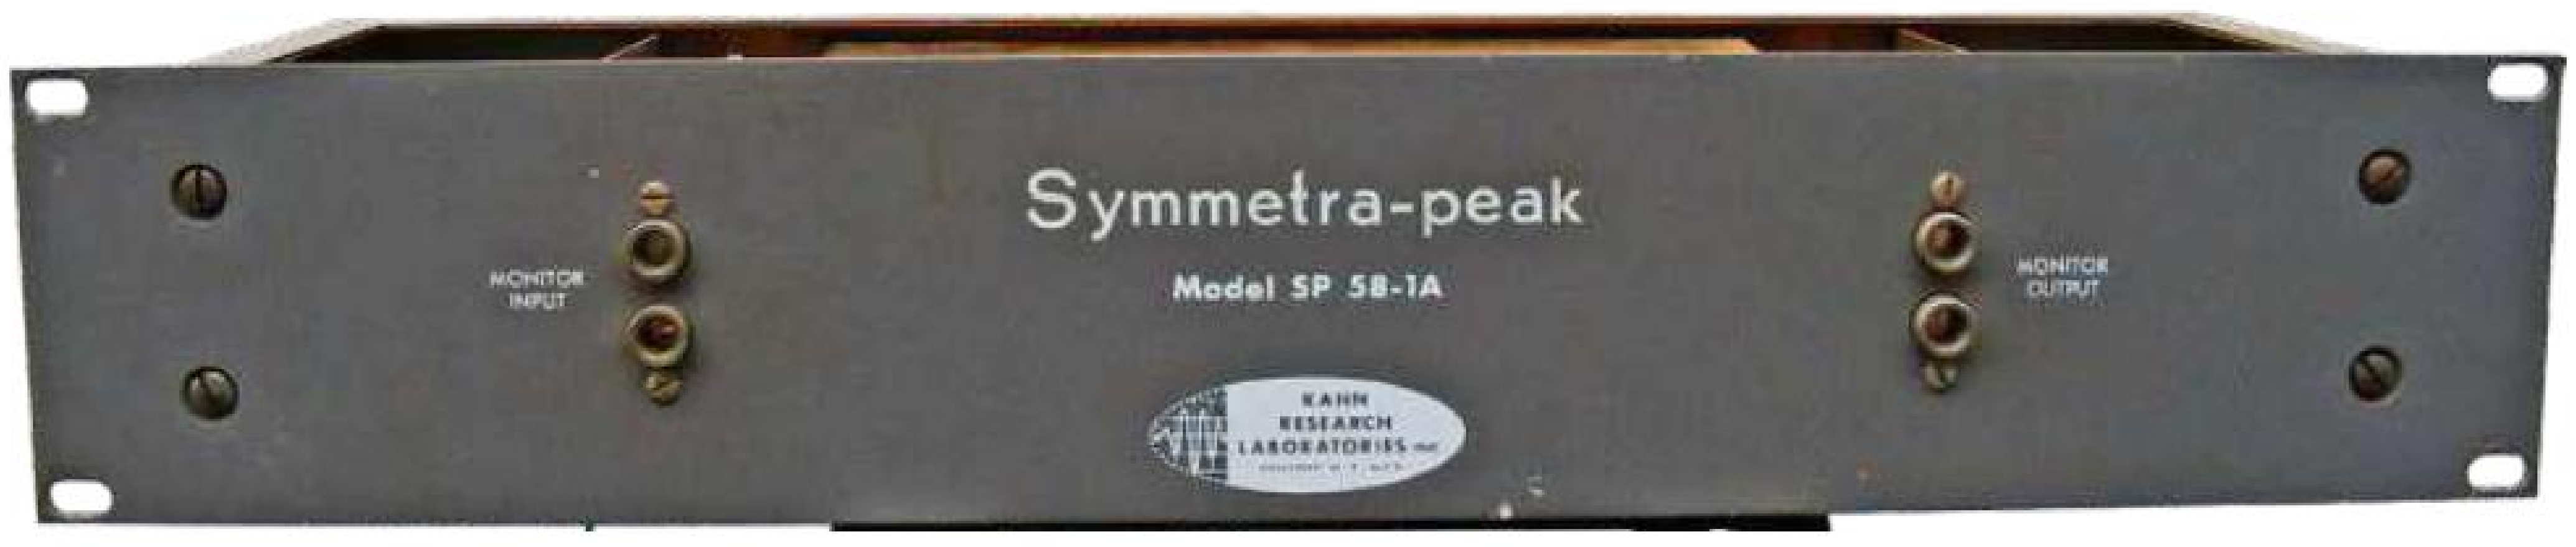

Photo 33: Symmetra-Peaks ensured maximum modulation of radio transmitters.

Symmetra-peak
Model SP 58-1A
MONITOR INPUT
MONITOR OUTPUT
KAHN RESEARCH LABORATORIES

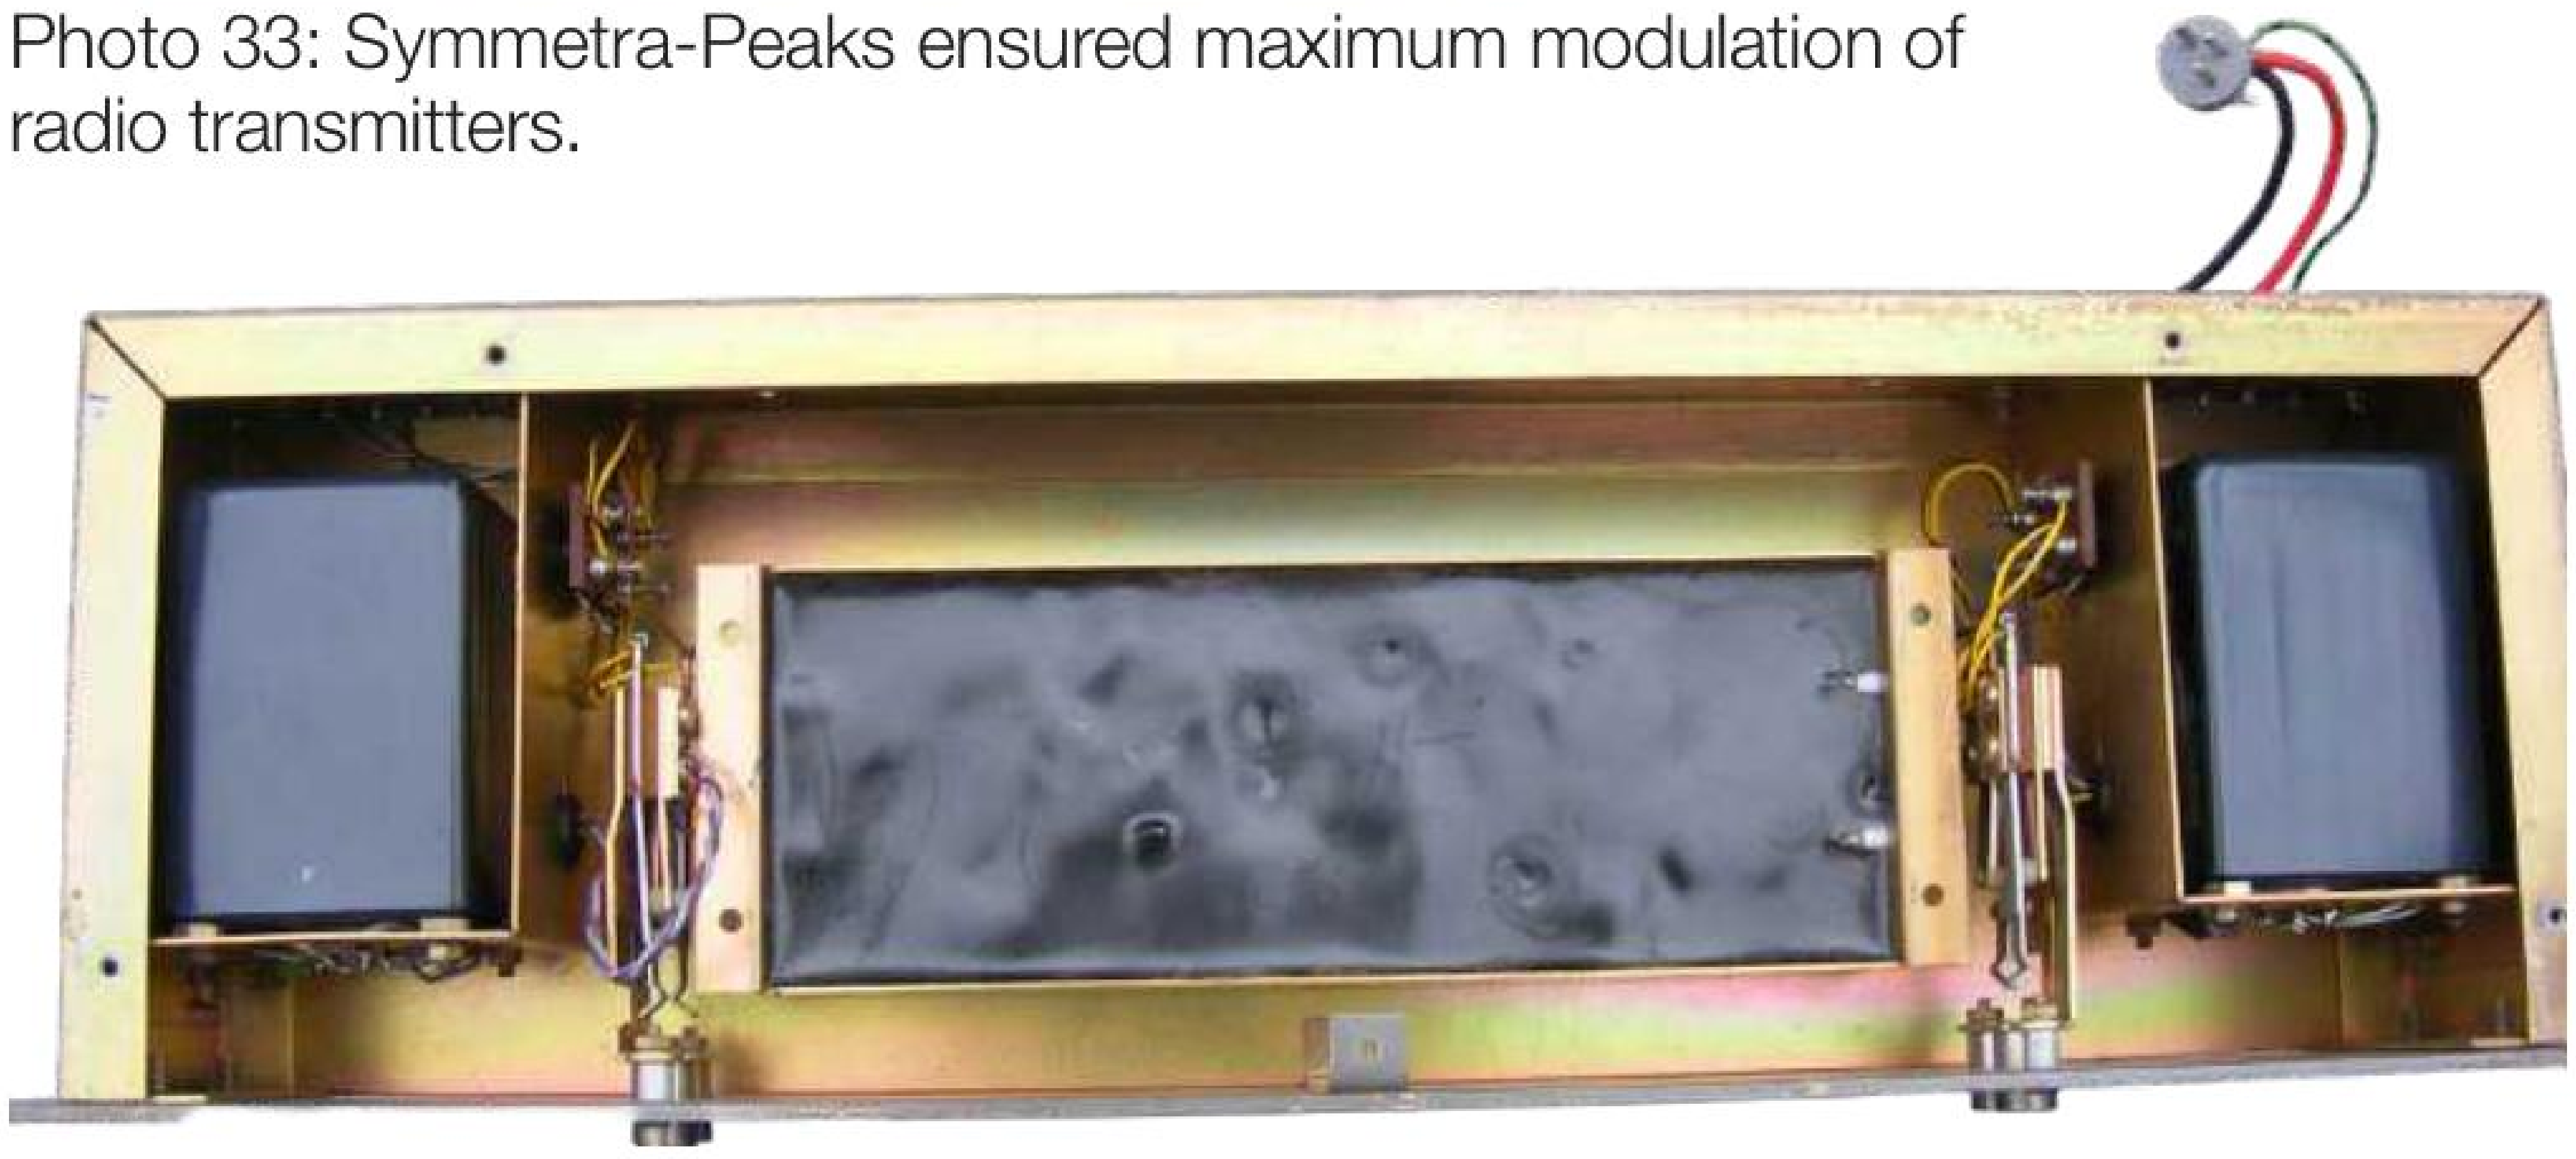

Photo 34: a full view of the Symmetra-Peak, showing the big audio transformers at each end and potted circuitry in between.

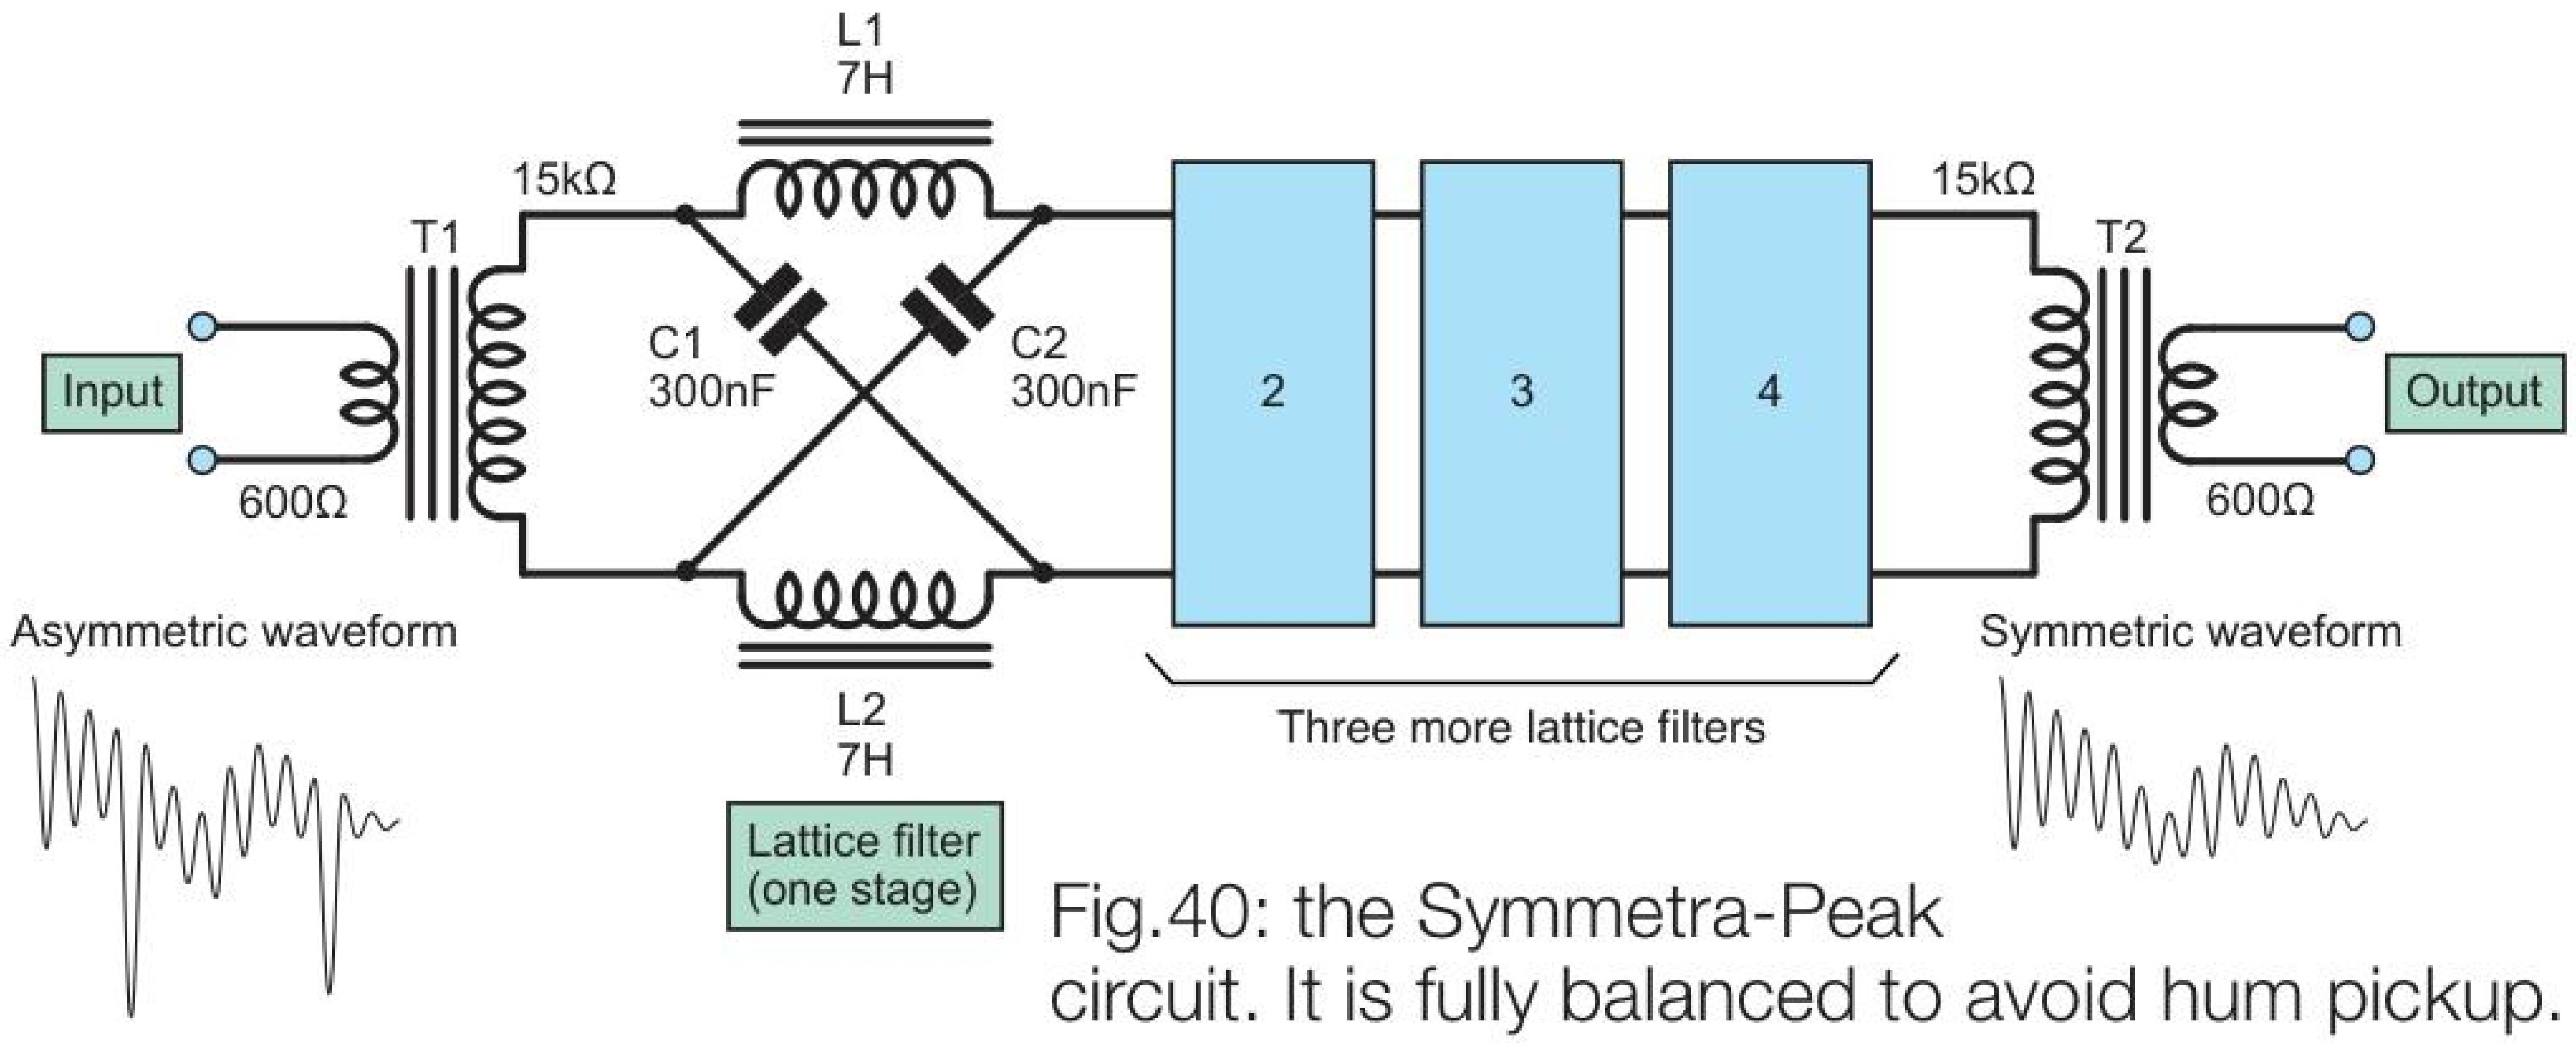

L1
7H
15kΩ
T1
C1
300nF
C2
300nF
2
3
4
15kΩ
T2
Input
600Ω
Output
600Ω
Asymmetric waveform
L2
7H
Lattice filter (one stage)
Three more lattice filters
Symmetric waveform

Fig.40: the Symmetra-Peak circuit. It is fully balanced to avoid hum pickup.

The most popular design was the Datong processor. This heterodyne-based technique is complex, but it pushes the intermodulation products out of the audio band.

All-pass filtering, along with high-speed amplitude compression, gives a clearer result, suitable even for studio work, for which this design is primarily intended. Compressors and clippers always work more consistently when fed with symmetrical waveforms, because DC shifts due to the asymmetry are avoided.

The centre frequency chosen is normally 200-400Hz, depending on the speaker's voice, to delay the fundamental relative to the harmonics. The number of poles required is at least eight. It is worth adjusting using one's own or a client's voice. Male voices are more asymmetric, lower in frequency, with more harmonics than female voices and thus respond better to this technique.

The question is: how to optimise the phase shift? It's all very well using an oscilloscope to check waveform symmetry on the bench, twiddling the pot while making various vocal sounds. It occurred to Harvey that a meter would be a good idea.

I thought that rectifying the upper and lower halves of the waveform separately, averaging them out, then reading the difference with a centre-zero meter might do the job. I'm working on it. Alternatively, there are some cheap LCD oscilloscope modules out there, such as the Kuman DSO 138 kit.

Broadcast symmetriciser

Now for a bit of history. Thanks to Harvey's extensive eBay searches, I'm able to show you a unique bit of gear called the Symmetra-Peak in Photo 33. As most people know, radio in the USA is still very important, and there has always been huge commercial competition between stations. This means that 'standing out' on the tuning dial, such as having the loudest sound, was crucial.

In 1959, Mr. Leonard. R. Kahn designed this secret processor box that was sufficiently Hi-Fi to symmetricise the signal going to a radio station transmitter, giving maximum modulation. It was an expensive bit of audio engineering, but the increased listenership and consequent increase in advertising revenue made it a worthwhile investment.

Inside was an epoxy-encapsulated passive lattice all-pass filter terminated by input and output balancing transformers, as shown in Photo 34. Luckily, I didn't have to dissolve the epoxy with toxic solvents to bring you the circuit, since it was already in the patent. This is shown in Fig.40.

Notice how the transformers have a high turns ratio.

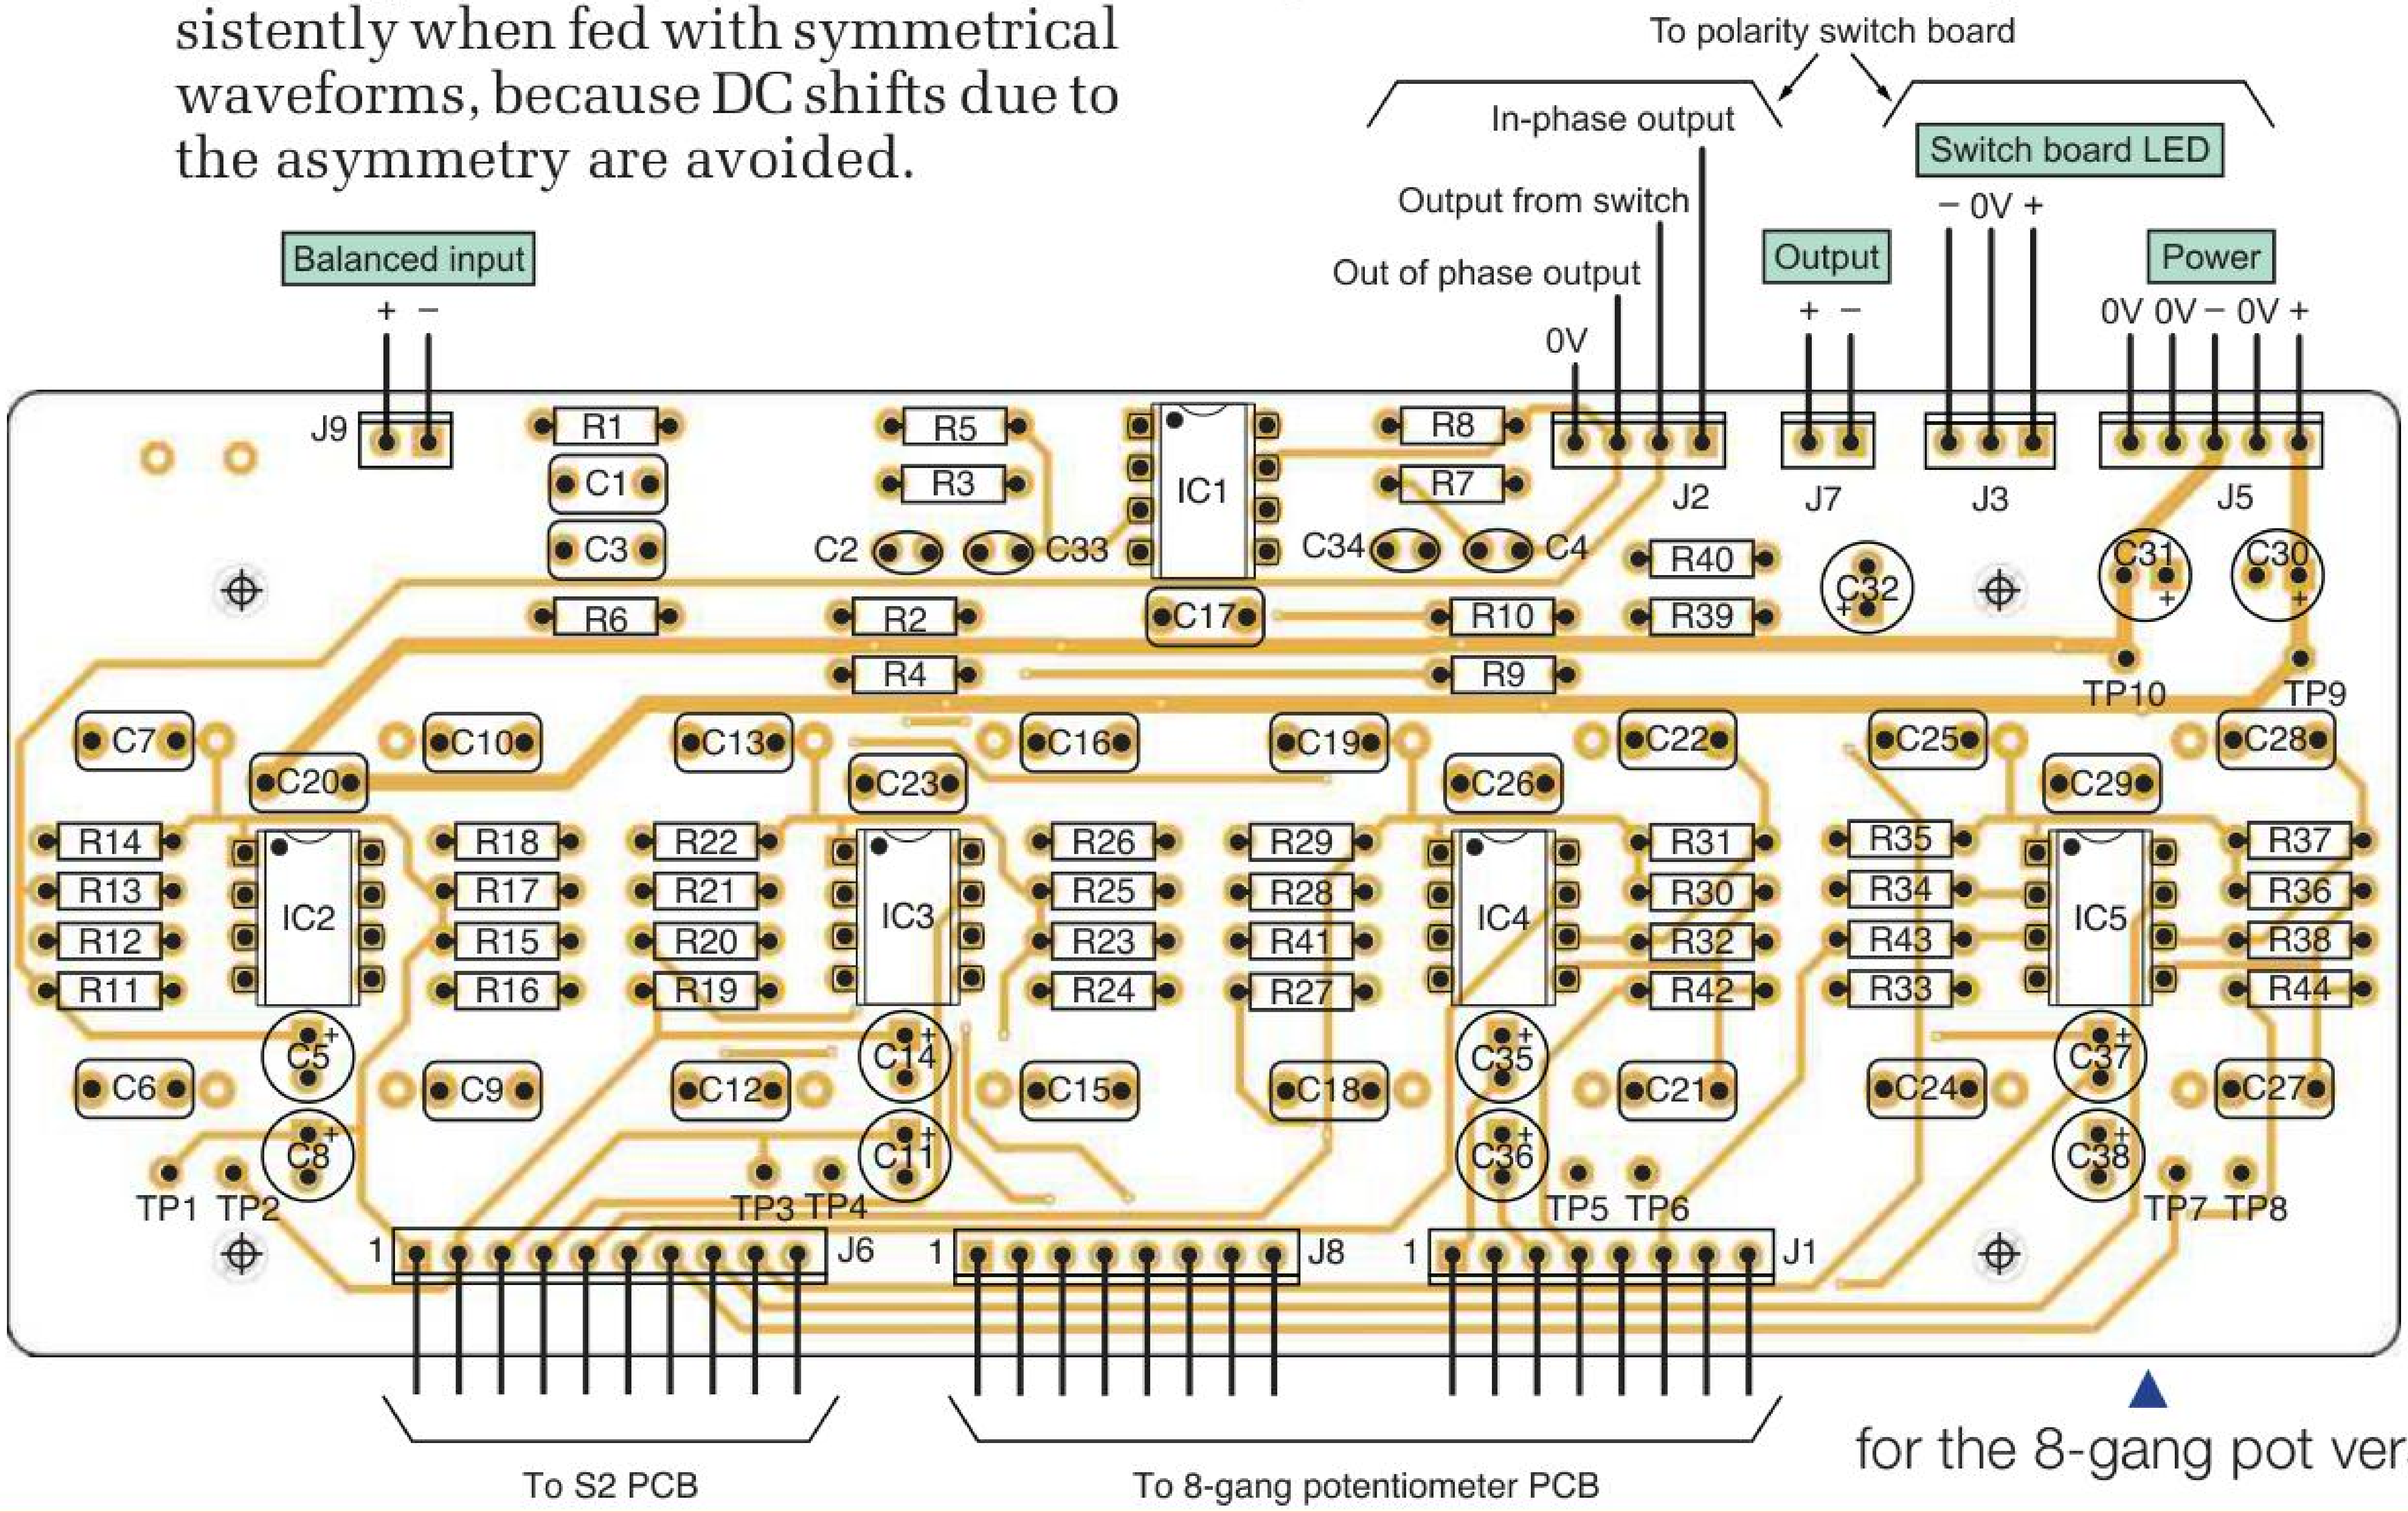

To polarity switch board
In-phase output
Switch board LED
− 0V +
Output from switch
Out of phase output
0V
Balanced input
+ −
Output
+ −
Power
0V 0V − 0V +
J9
J2
J7
J3
J5
R1
C1
C3
R6
R5
R3
C2
C33
R2
R4
IC1
C17
C34
C4
R8
R7
R40
R10
R39
R9
C32
C31
C30
TP10
TP9
C7
C20
C10
C13
C23
C16
C19
C26
C22
C25
C29
C28
R14
R13
R12
R11
IC2
R18
R17
R15
R16
R22
R21
R20
R19
IC3
R26
R25
R23
R24
R29
R28
R41
R27
IC4
R31
R30
R32
R42
R35
R34
R43
R33
IC5
R37
R36
R38
R44
C5
C14
C35
C37
C6
C9
C12
C15
C18
C21
C24
C27
C8
C11
C36
C38
TP1 TP2
TP3 TP4
TP5 TP6
TP7 TP8
1
J6
J8
J1
To S2 PCB
To 8-gang potentiometer PCB

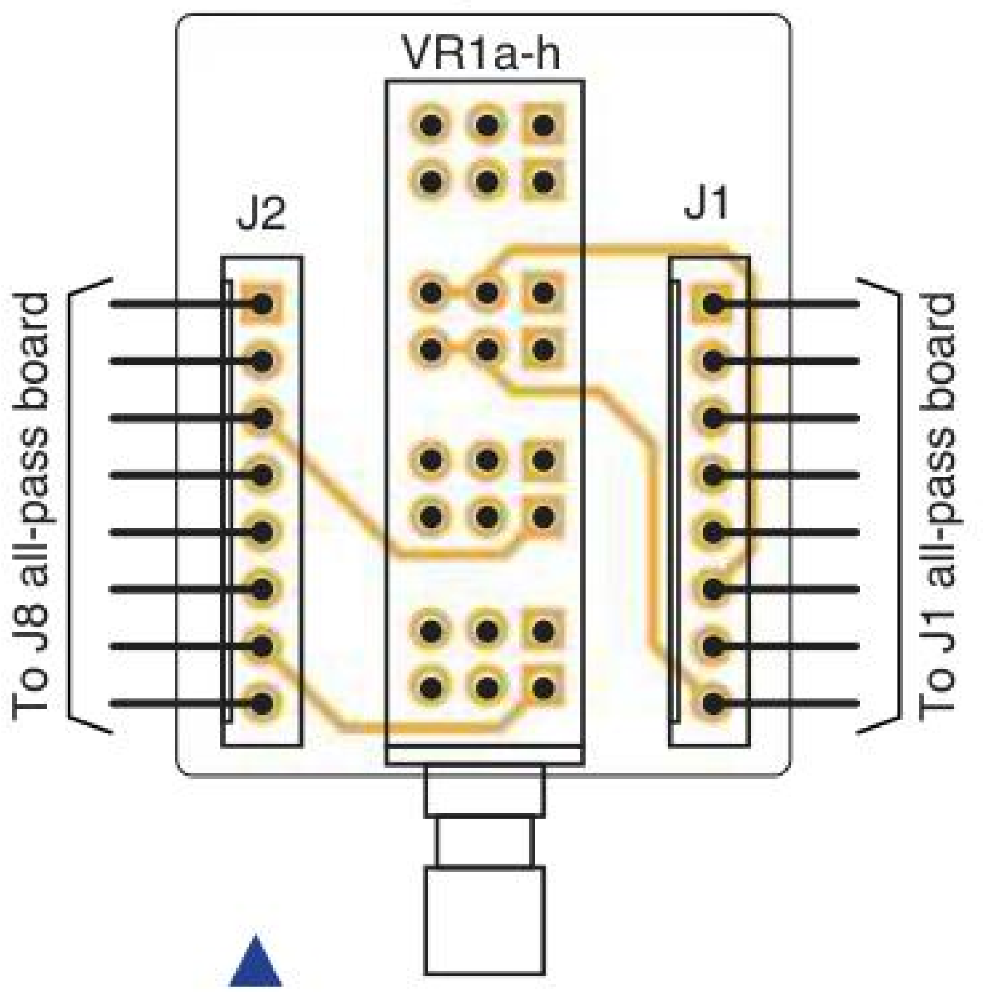

VR1a-h
J2
J1
To J8 all-pass board
To J1 all-pass board

Fig.39: the PCB overlay for the separate 8-gang pot board.

Fig.38: the PCB overlay for the 8-gang pot version of the all-pass filter board.

This was necessary to match the high impedance (15kΩ) of the lattice filter to the low 600Ω line impedance that was the audio standard of the time.

The filter was high impedance because the only passive audio components available in the valve era were high-impedance inductors and relatively low-value capacitors, with 7H and 30nF being used. If you were to attempt this passive approach today, modern pot cores of a few 100s of millihenries and high-value polypropylene capacitors of several µF would be used, with no transformers.

A high-current buffer stage would be needed to drive it, though. Audio inductors are expensive, so the more cost-effective way is to use active op amp circuitry to eliminate them altogether (eg, using gyrators). The lattice can be replaced with four op amp all-pass filter stages.

Mr. Kahn found that a good compromise between complexity and performance was to make the delay around 5ms below 300Hz and under 3ms above 300Hz. The effect of the Khan Symmetra-Peak on a square wave is shown in Photo 35. This reduced the energy deviation on general audio signals with an energy imbalance of 20-30% down to 1-2%.

The phase plot of the Symmetra-Peak is shown in Fig.41. Note that each vertical line on such plots denotes a polarity reversal of 360°, each representing two poles of 180°. Therefore, the unit has four poles. 360° is often called a phase rotation; all-pass filters are sometimes called phase rotators.

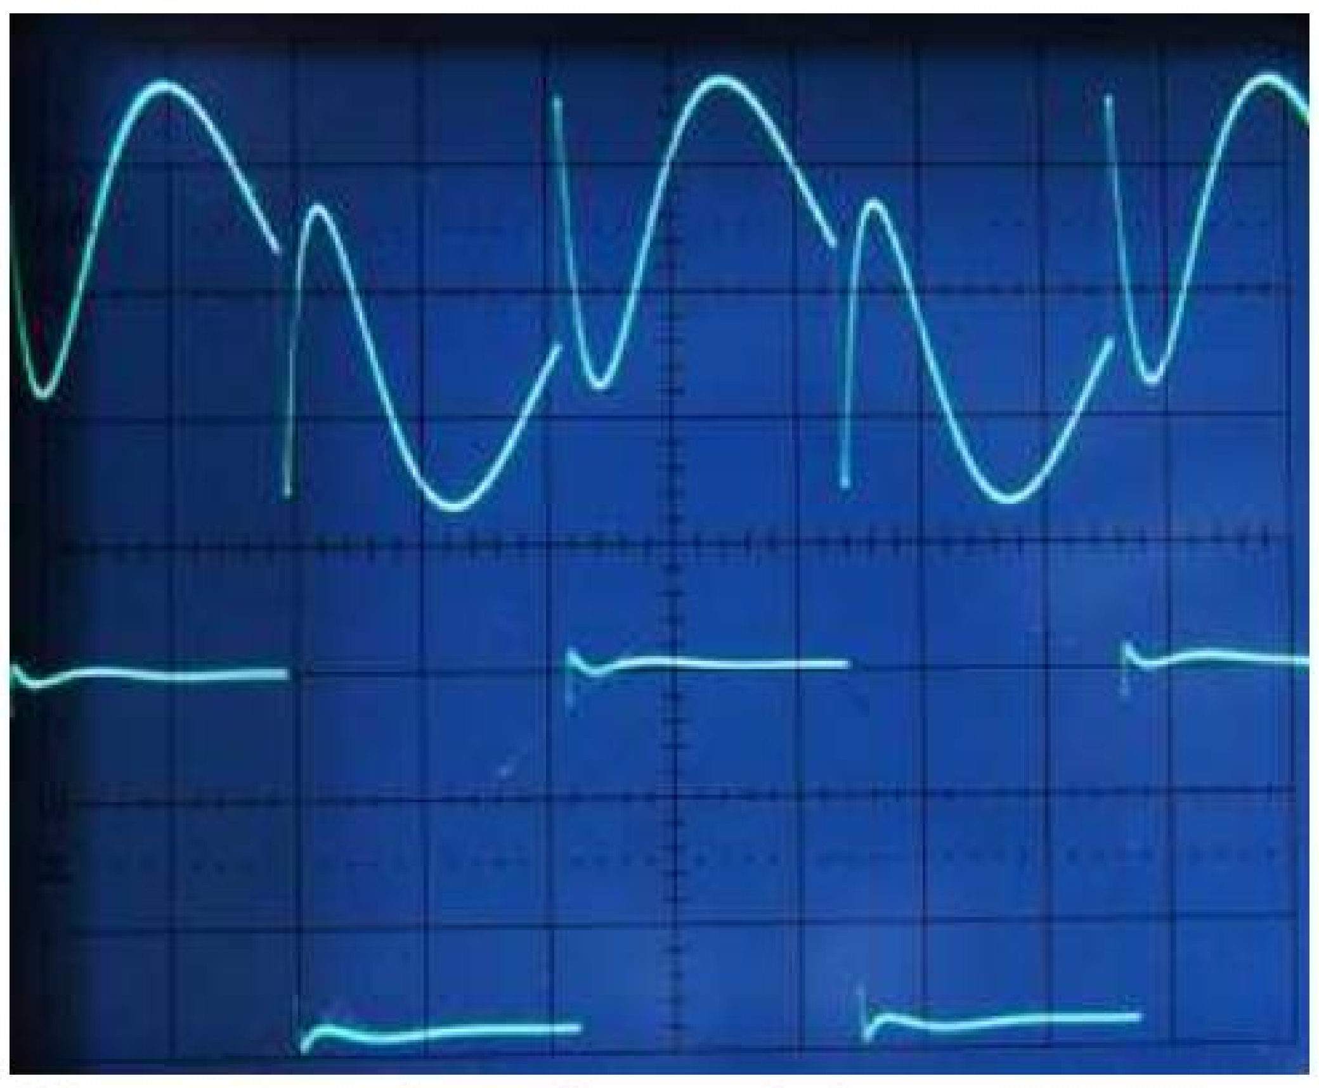
Photo 35: the effect of the Symmetra-Peak on a square wave.

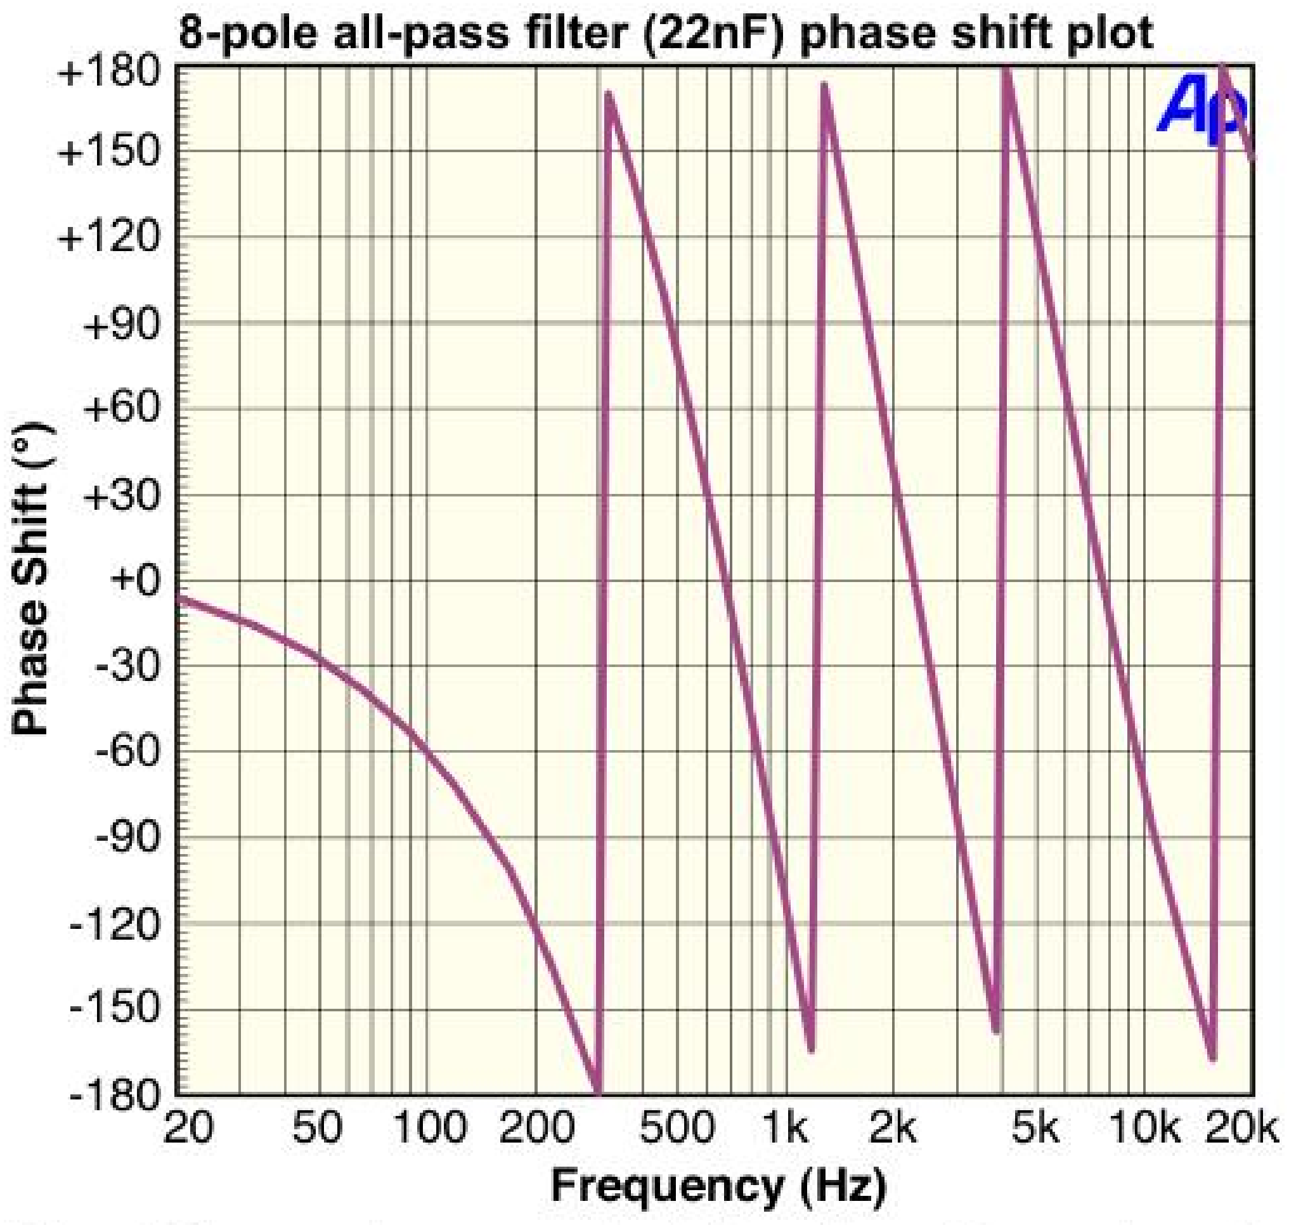

Fig.42: a phase plot for the 8-pole all-pass filter with 22nF filter capacitors. The pot is in the middle position.

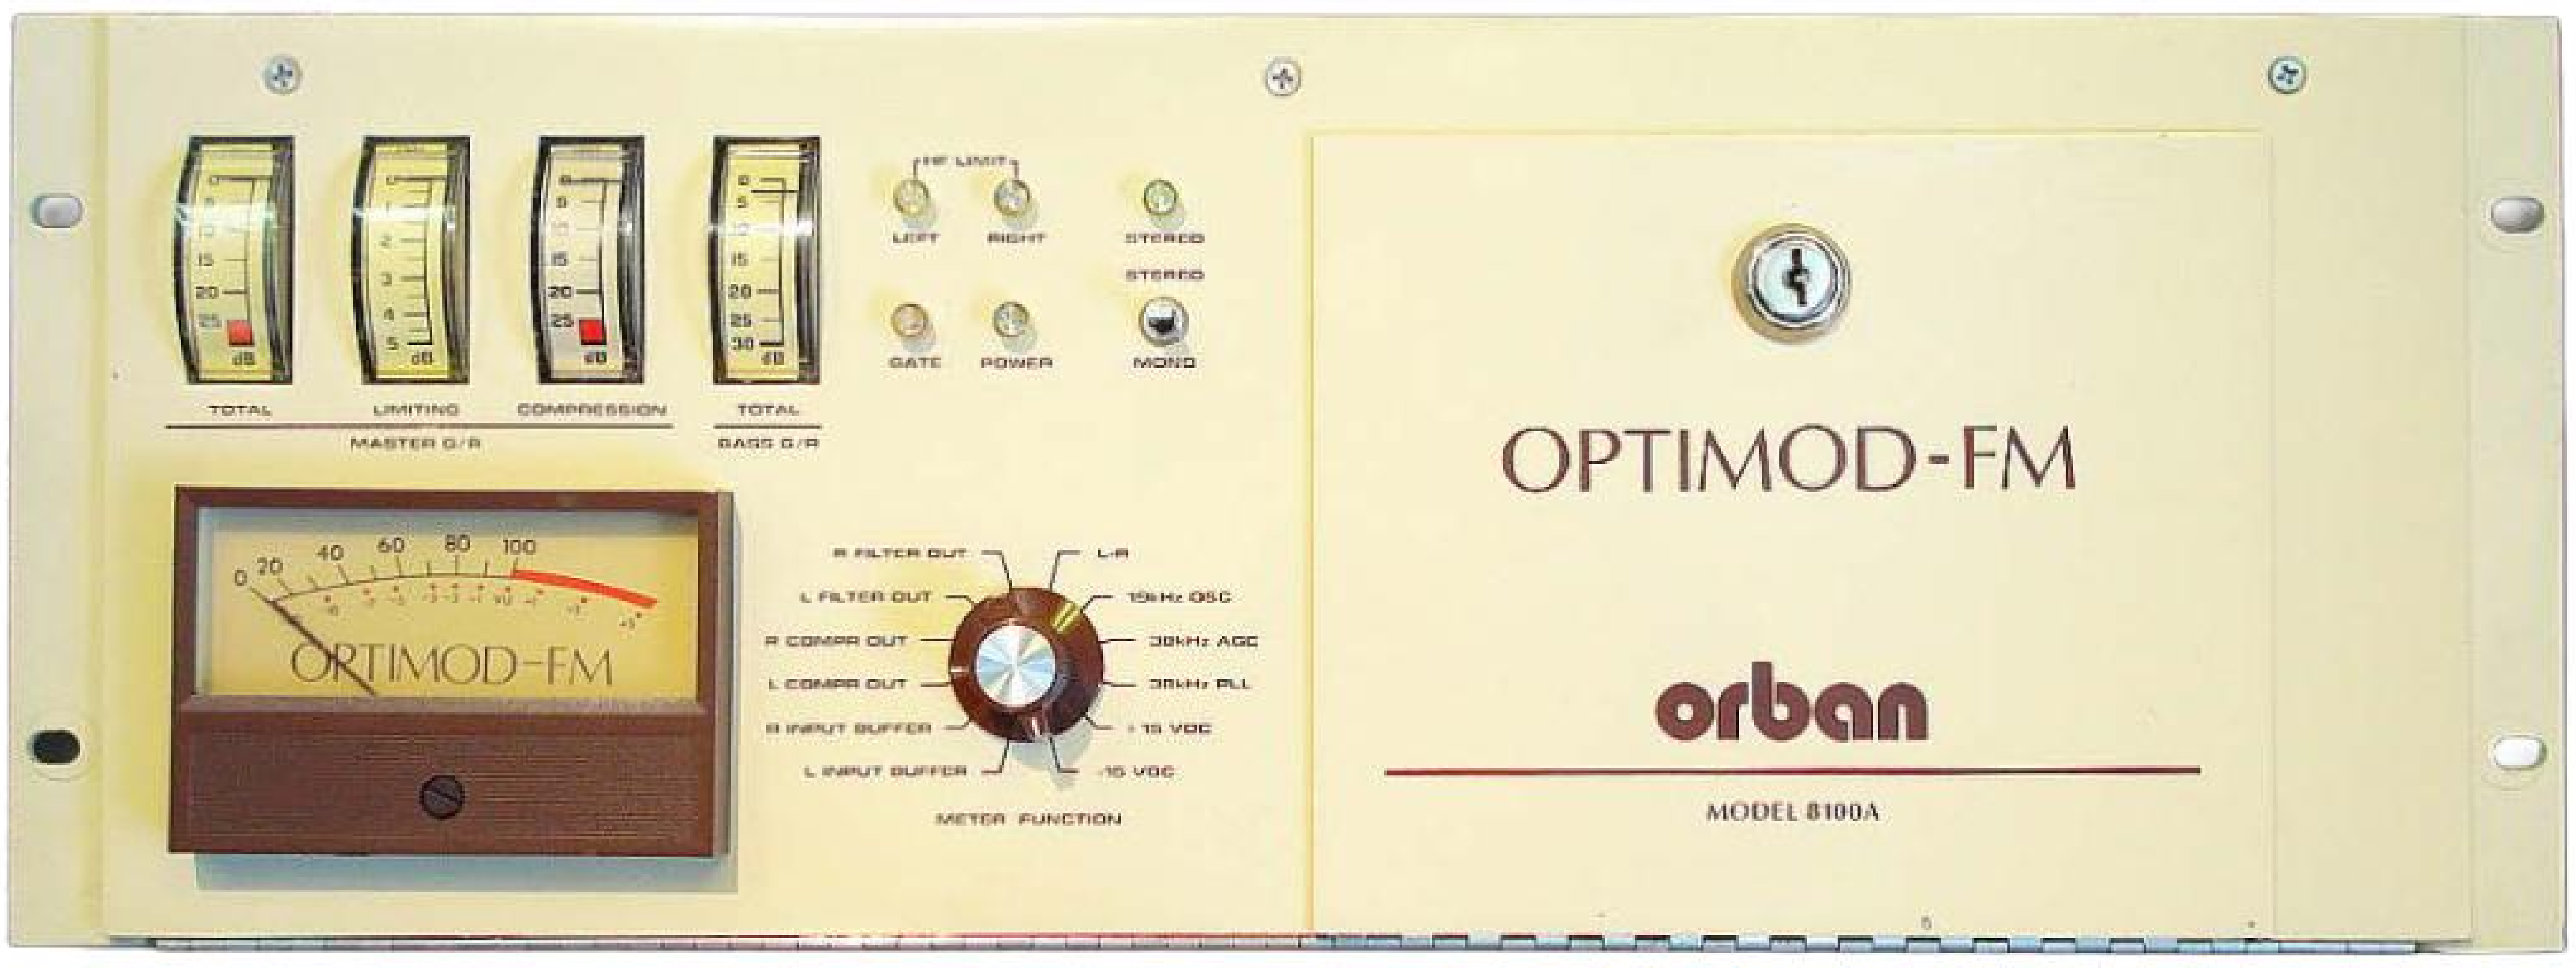
OPTIMOD-FM
orban
MODEL 8100A
Photo 36: the Orban Optimod. Source: www.proaudio.com

This plot was a pig to make. I had to get Harvey to set up the Audio Precision analyser software.

There have been a few all-pass speech symmetricisers produced for the ham radio scene. Research by Gary Blau (amateur radio callsign and website W3AM) and by James L Tonne (callsign W4ENE) has resulted in an active version of the Symmetra-Peak with the frequency set to 340Hz (R = 10kΩ and C = 47nF).

They doubled the number of poles from Kahn's original four to eight, which further improved the symmetry. With a 50kΩ anti-log Bourns pot and 22nF capacitors, this frequency is obtained with the pot at around one-third rotation, as shown in Fig.42.

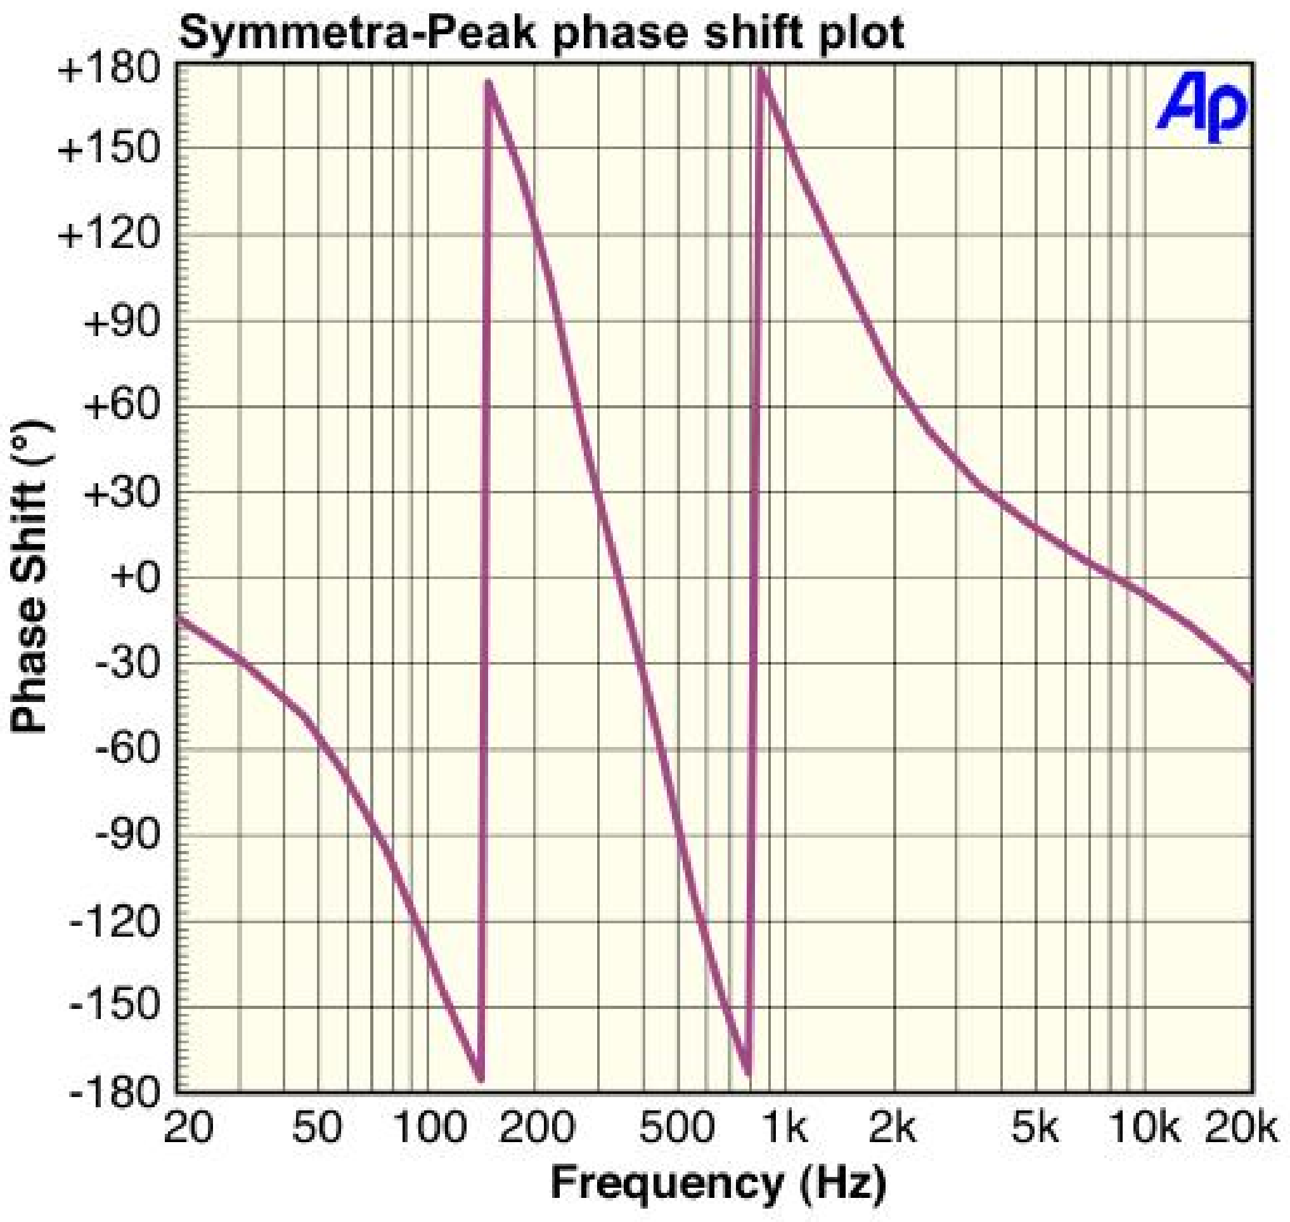

Fig.41: the Symmetra-Peak phase plot. This indicates that it uses a four-pole all-pass filter.

The Optimod

The Optimod started a trend for units to optimise the modulation of transmitters. It culminated in the Orban Optimod, a five-band multi-way limiter. The internal active crossover provides a large amount of phase shuffling, along with a bit of dedicated HF all-pass delay, to provide ultra-fast 'look ahead' limiting of peaks.

It was so effective that artists said their records sounded better on the radio than in the studio!

There were two versions: a mono version optimised for AM transmission, and a larger stereo one for FM, shown in Photo 36. I saw an AM unit at Retrotech last year for £300 and regretfully didn't buy or photograph it. They were originally very expensive, built to almost military standards, with hermetic metal-can op amps.

Barcus Berry Sonic Maximiser

Attempts were made to make similar processors more affordable by companies such as Aphex and Circuit Research Labs, but the biggest seller was the BBE Sonic Maximiser.

This was a summed second-order output state variable filter, which provides some all-pass action in

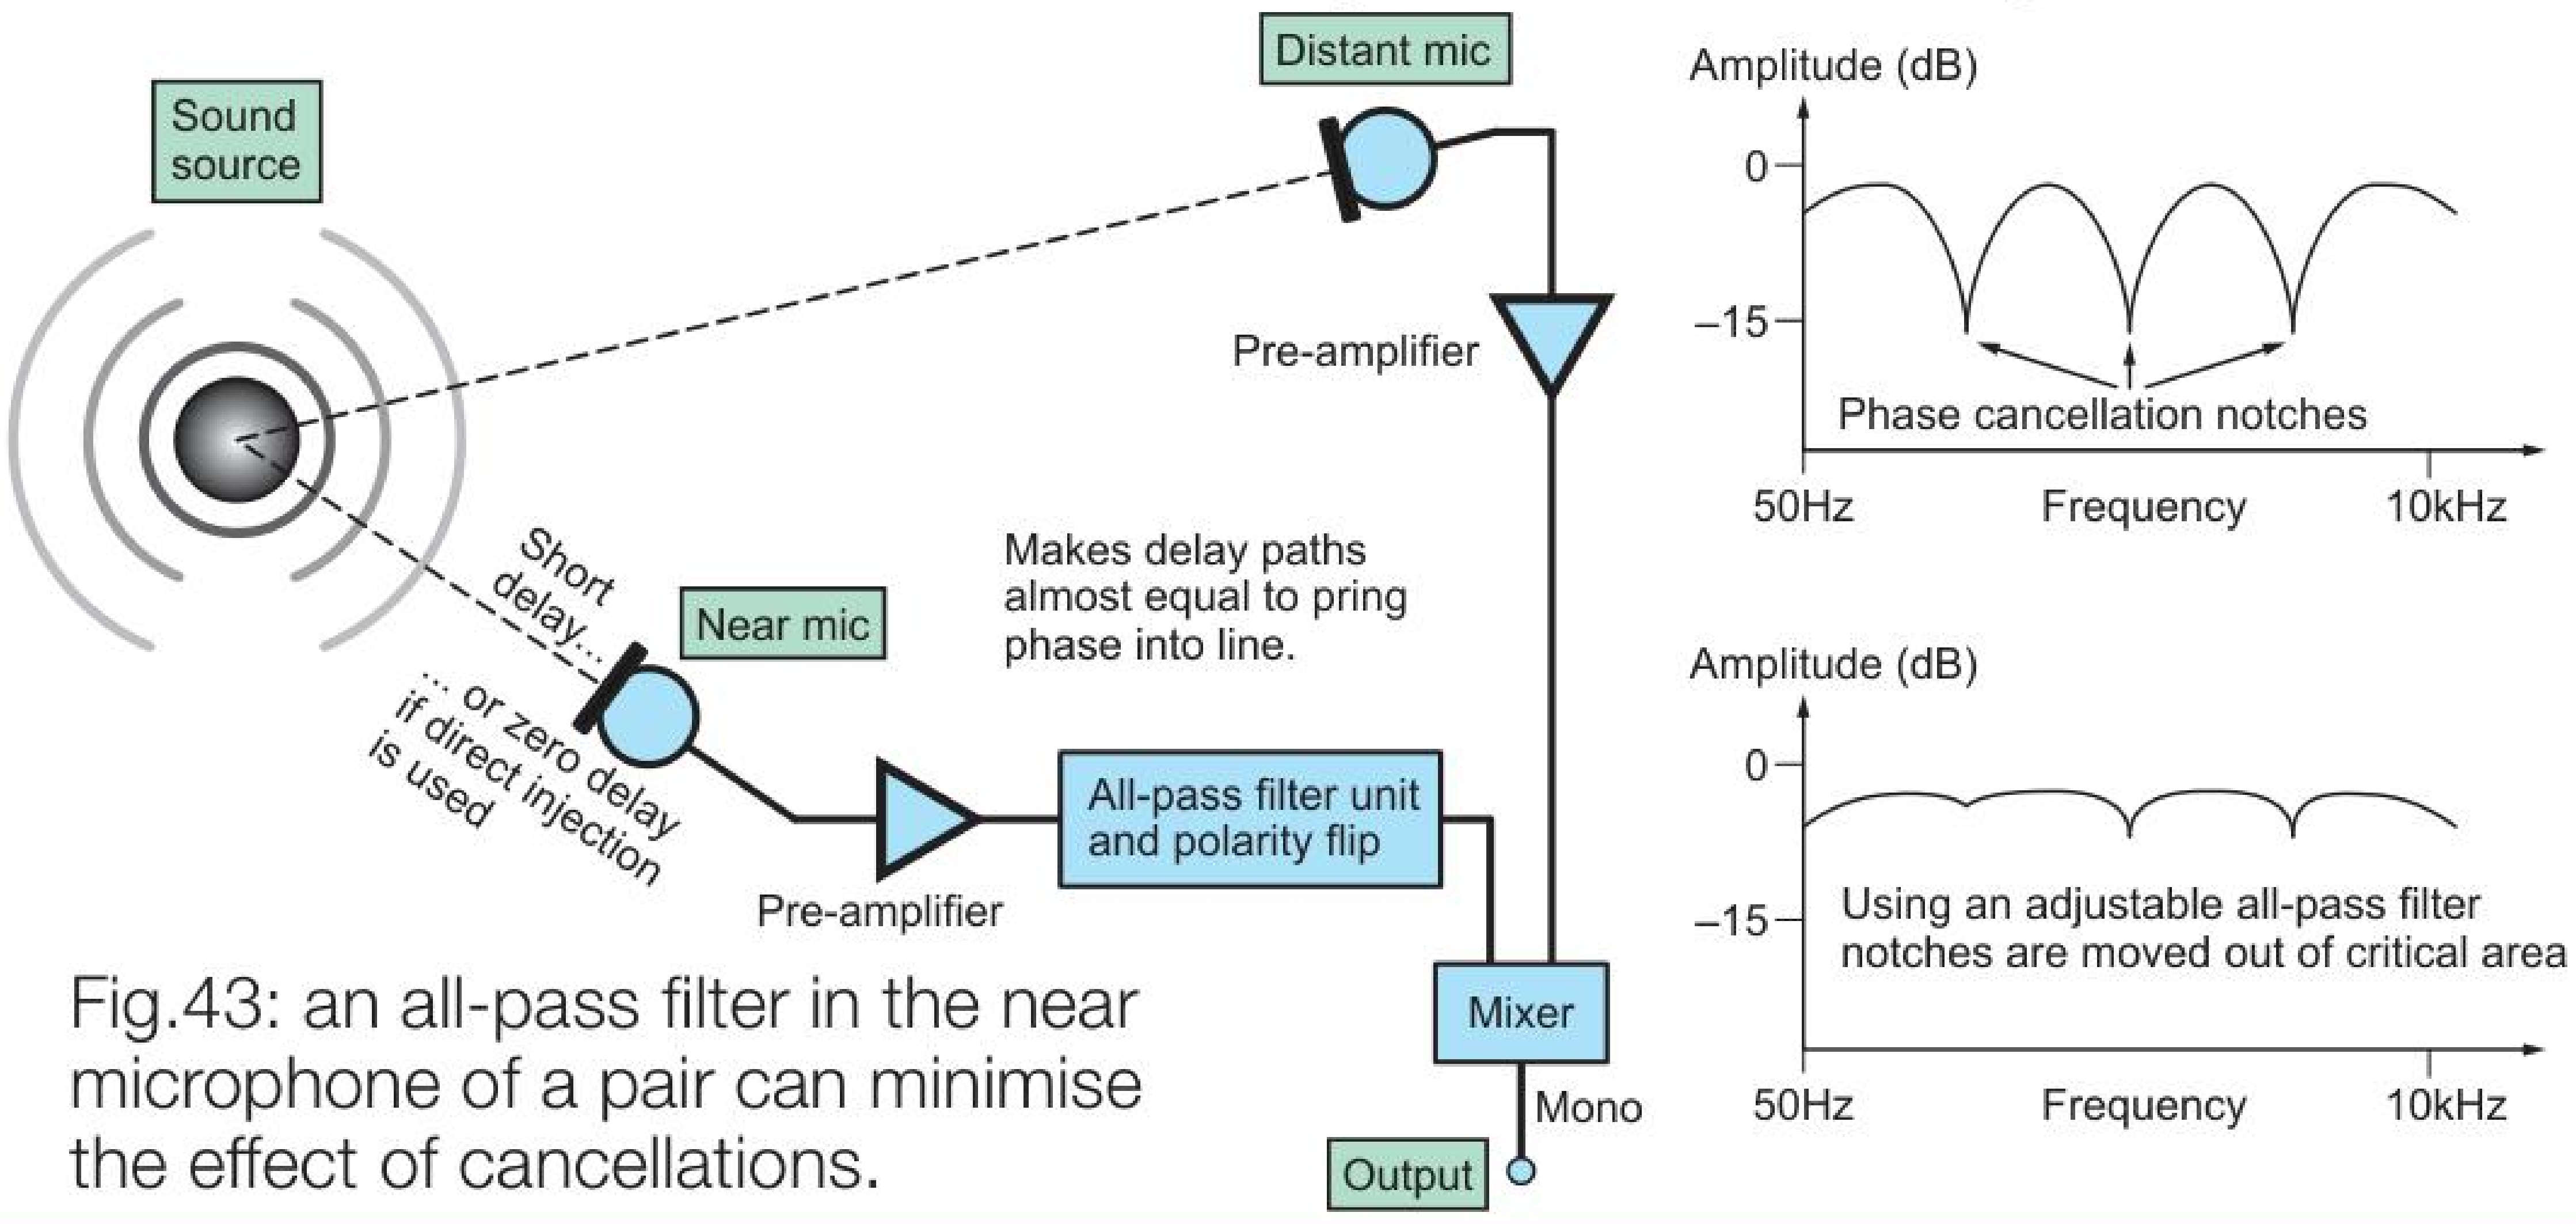
Sound source
Distant mic
Pre-amplifier
Short delay...
... or zero delay if direct injection is used
Near mic
Makes delay paths almost equal to pring phase into line.
All-pass filter unit and polarity flip
Pre-amplifier
Mixer
Mono
Output
Amplitude (dB)
0
–15
Phase cancellation notches
50Hz
Frequency
10kHz
Amplitude (dB)
0
–15
Using an adjustable all-pass filter notches are moved out of critical area
50Hz
Frequency
10kHz
Fig.43: an all-pass filter in the near microphone of a pair can minimise the effect of cancellations.

PHASE
INVERT
SHIFT
ON
180°–360°
0°
180°
FILTER
RANGE
CUT-OFF
ON
300Hz–3.8kHz
3kHz
300Hz
38kHz
3.8kHz
Radial engineering
GROUND
POWER
PUSH TO LIFT

Photo 37: the Radial Engineering Phazer, available from Thomann.

instrument in
Little Labs
ibp analog phase alignment tool
line/ instrument
earth lift
phase adj bypass
phase invert
phase adj 90° 180°
phase center lo/hi
phase adjust

Photo 38: Little Labs. Source: eBay.

the normal way, but it is little more than a bass/treble booster circuit with steeper slopes, like a Hi-Fi Baxandall tone control on steroids.

There is a dynamic noise-filtering VCA circuit on the treble output, which avoids the noise-boosting problem that aural exciters typically suffer from. Of course, the marketing hype has the usual exhortations about psychoacoustic phase correction, reinserting lost information etc, shrouding what is, in fact, quite basic audio circuitry in secrecy.

The circuit was disclosed in the data sheet for a BBE chip, the NJM2153, made by JRC. The unit worked well on the low-resolution sampled drum beats of the early 1990s. Sadly, it didn't work on smooth atmospheric or orchestral music.

The use and abuse of phase

Because of the Optimod, artists wanted that sound in the studio, and units such as the Aphex Dominator were made to achieve this. The result was that the music was subjected to 'Optimod style' processing twice if played on the radio. This didn't sound twice as good; the result was a strange, loud audio 'mush'.

The Red Hot Chili Peppers and their 1999 album *Californication* took this to extremes; it is a loudness landmark. It was so heavily processed that the waveform contains flat clipped sections within the waveform, not merely at the peaks. This is what can happen if you phase-shuffle audio that has already been clipped!

It sounds unnaturally loud, which is great for grabbing attention in adverts, but fatiguing for music listeners, especially over a whole album. Thankfully, the loudness wars seem to have peaked in the mid-2000s, no pun intended.

[Editor's note: the loudness war was so bad that Rush's 2002 album Vapor Trails *was remixed and re-released in 2013 after years of complaints about its harsh, fatiguing sound. Even the band members agreed! Metallica's 2008 album* Death Magnetic *was another infamous example; later digital releases used a less punishing master.]*

Loudness standards such as the European Broadcasting Union's R 128, along with loudness normalisation on streaming services, finally made this kind of nonsense far less common.

Phase alignment unit

An all-pass filter is possibly the last effects unit a studio would buy; phase alignment units are very expensive, typically £350 due to the low demand, so it's a box worth making yourself. They are mainly used in the studio to ensure that a sound source recorded with multiple microphones or a directly injected instrument, such as a bass guitar, with a mic'd-up amp speaker, sum properly.

This trick makes it possible to avoid frequency notches caused by cancellation in critical frequency areas, as shown in Fig.43. Of the few units made, the only ones I've used are the Radial Engineering Phazer in Photo 37 and the Little Labs Analogue Phase Alignment Tool (Photo 38) which they have at my local studio, Giant Wafer.

For the eight-gang pot all-pass filter, the total amount of available phase shift is much greater than these units. The phase shift at opposite ends of the pot rotation is shown in Figs.44 & 45.

An unfortunate news flash

We'll have to interrupt the phase project now, since I've been informed by former PE contributor, Ben Duncan, of a looming audio industry chip supply problem. Ben Duncan introduced me and many home constructors to the original Philips TDA1034, later becoming the Signetics NE5534 in September 1978. The dual version, the NE5532, is now ubiquitous.

These original devices are shown in Photo 39. However, the analog audio world is now about to suffer a semiconductor price-gouging catastrophe. The full-specification NE5532 op amp is soon to be no more.

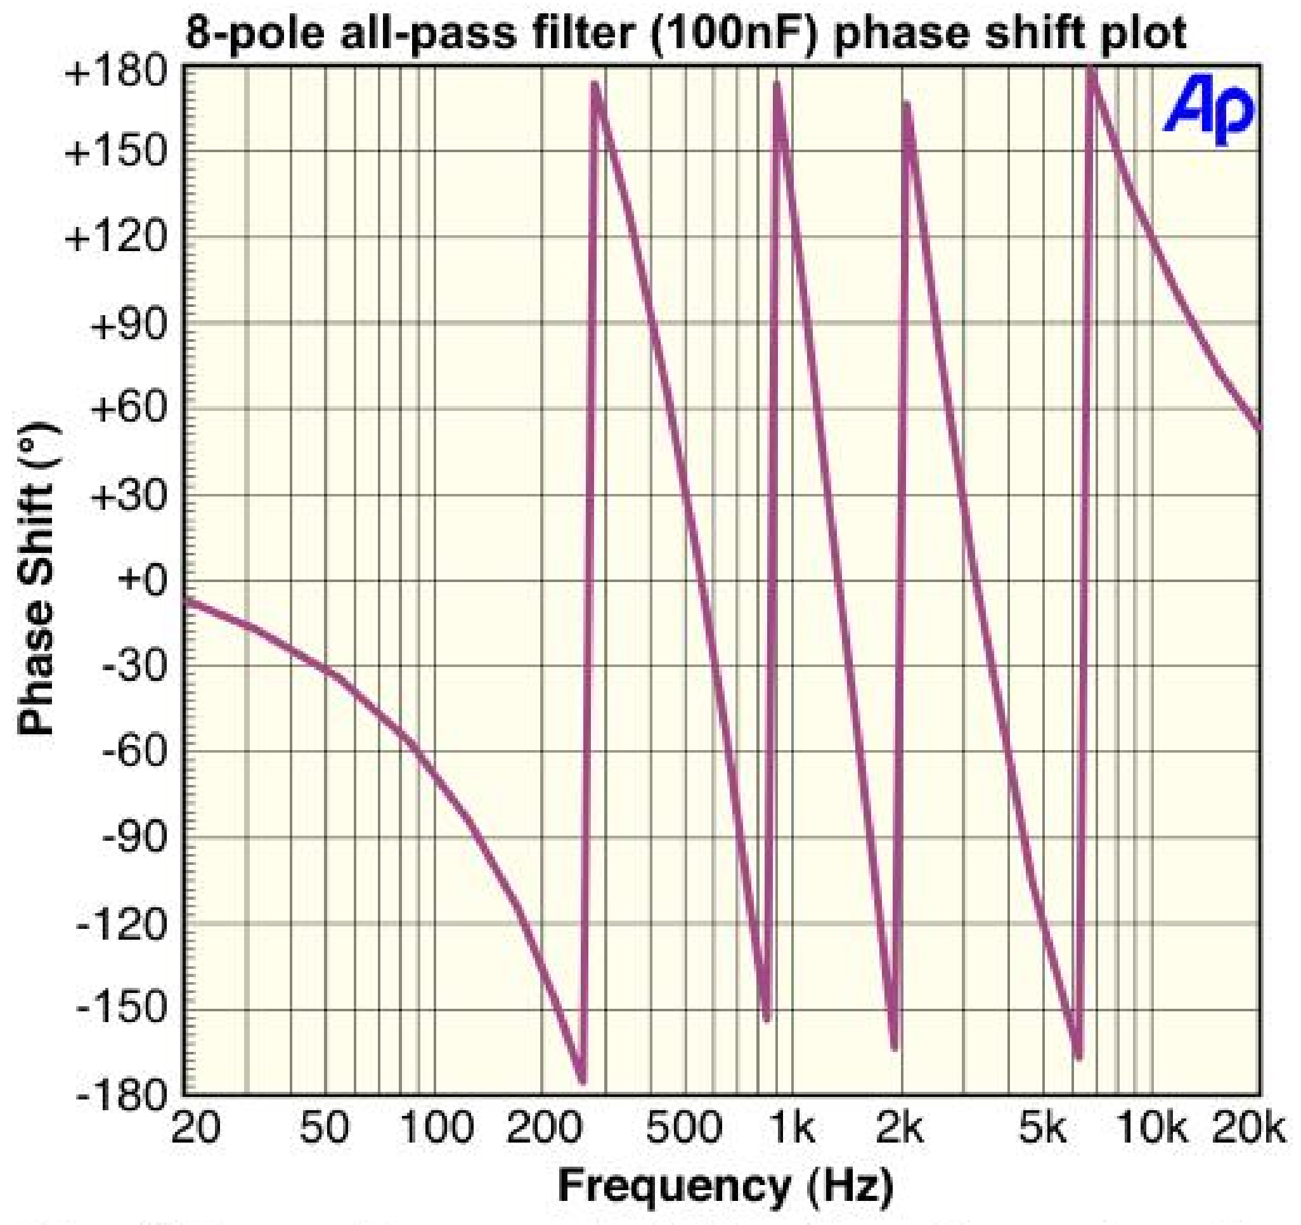

Fig.44: a phase plot for the 8-pole all-pass filter using 100nF caps with the pot fully clockwise for maximum delay.

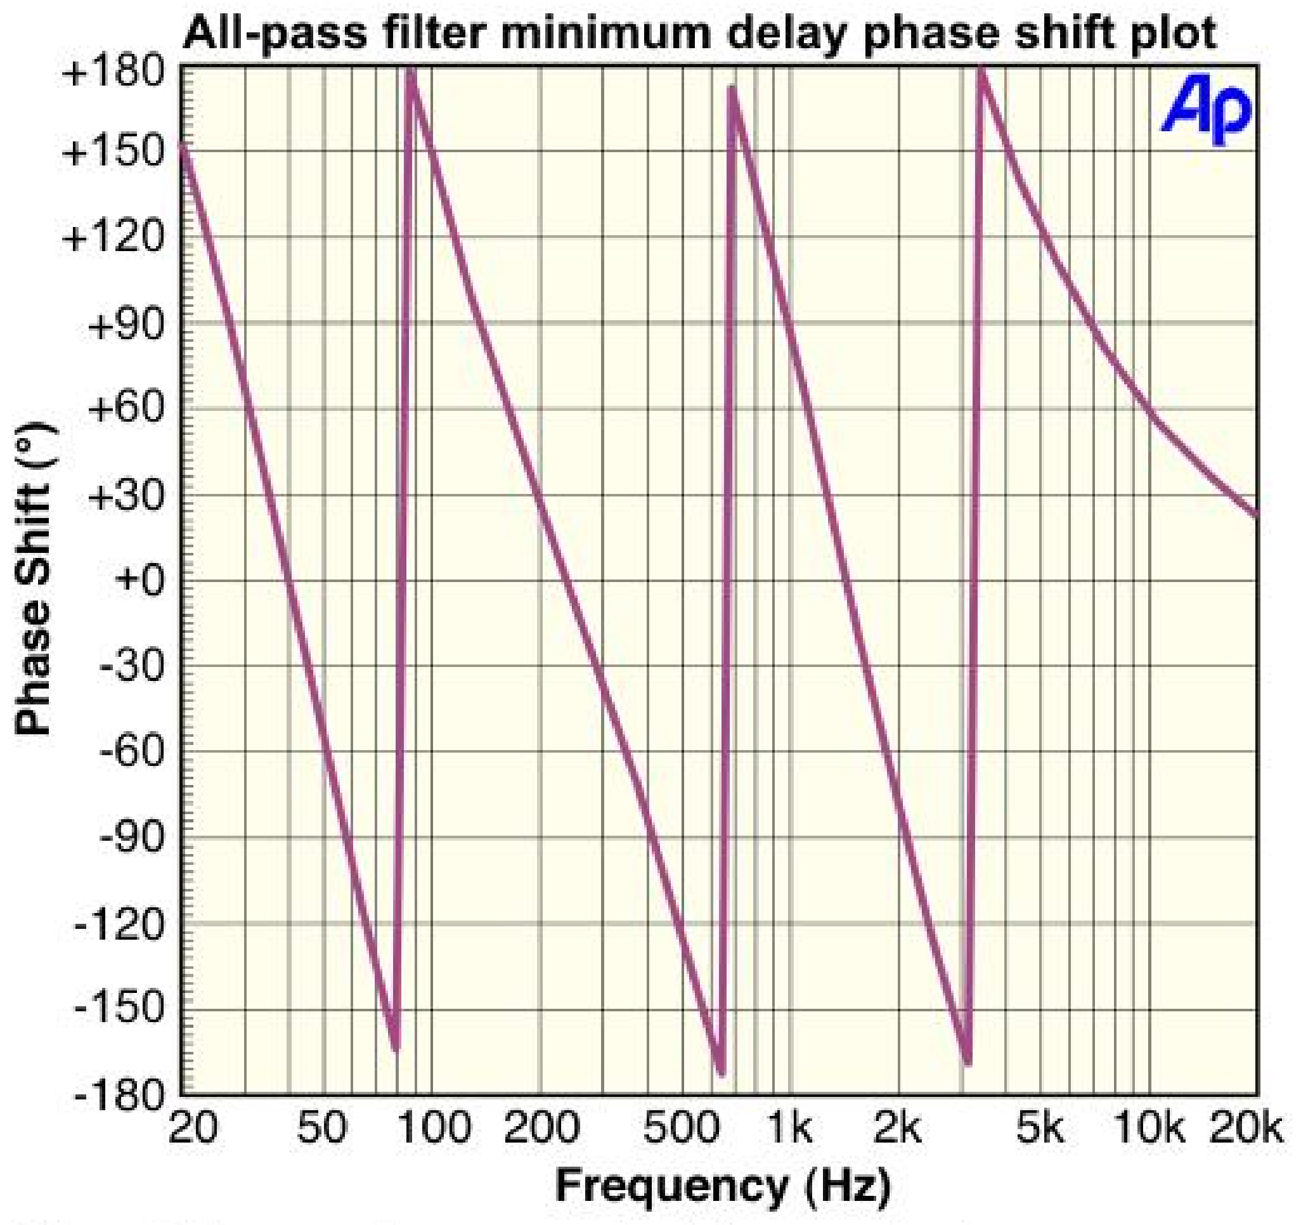

Fig.45: a phase plot for minimum delay, with the pot fully anti-clockwise.

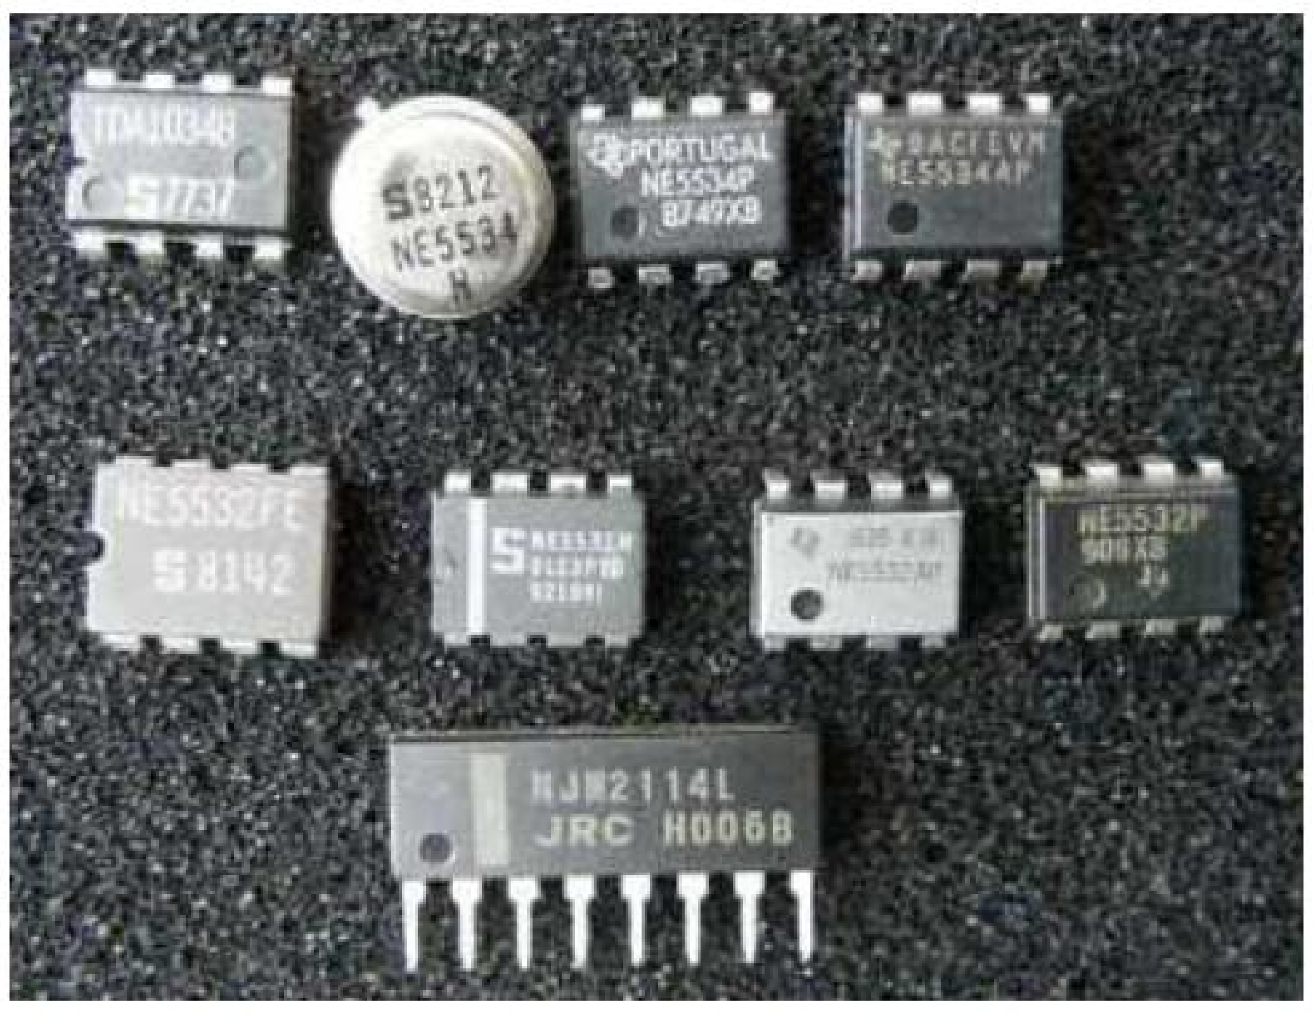

Photo 39: left-to-right, top-to-bottom: an original 1977 TDA1034, Signetics NE5534A, Signetics ceramic NE5532, various TI NE5532s and JRC's SIL NJM2114.

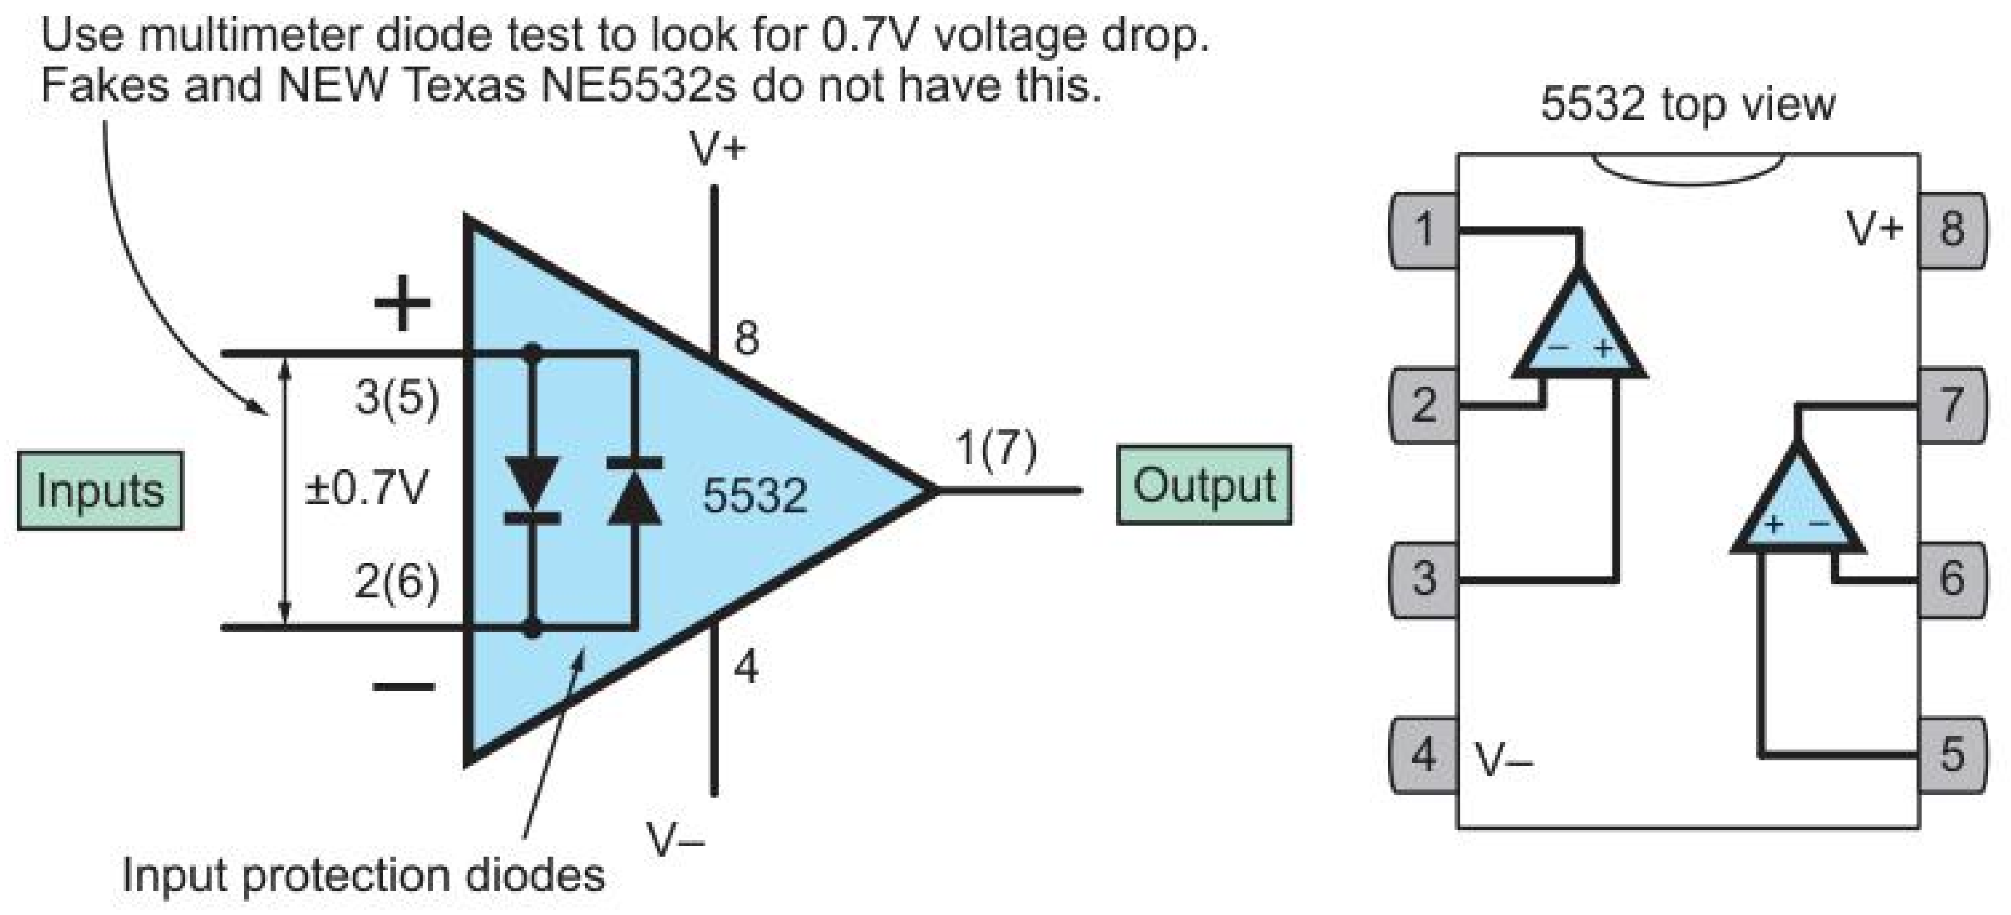
Use multimeter diode test to look for 0.7V voltage drop. Fakes and NEW Texas NE5532s do not have this.

Fig.46: checking for the internal diodes in an "NE5532" op amp IC.

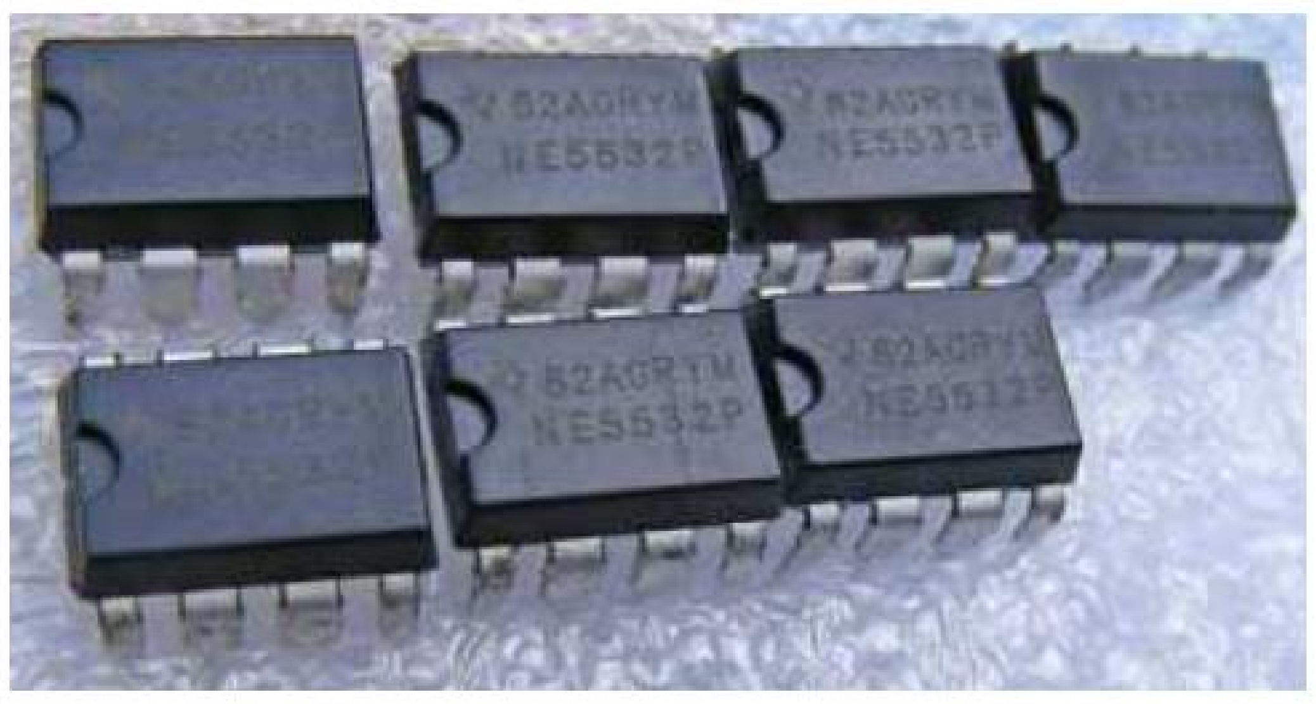

Photo 40: fake NE5532s from AliExpress. A bad copy or a completely different chip?

Dr Craig Sawyer, an associate of Ben's, has questioned Texas Instruments, the main producer, about changes in the specification. TI deliberately engineered the NE5532 and NE5534 down in February 2024, removing 8V from the maximum supply rail voltage, from ±22V to just ±18V!

The slew rate has also been reduced from 9/µs to 5/µs and the current consumption has dropped from 8mA to 6mA. The die geometry must have been shrunk as well, because the static voltage maximum has been reduced from 2kV to 1kV. One has to wonder if these are really the same parts, given their wildly different parameters from the original.

Was the NE5532 too good?

The original NE5532 was just too good for a low-priced generic product, and its wide availability was stopping customers from buying expensive, proprietary audio op amps. It's the usual accountancy-led dodge of die shrinkage and making legacy analog chips fit in with modern digital processes.

TI was originally castigated for making the high-frequency distortion of the NE5532 worse on their first versions compared to those made by Signetics when they took the company over. Also, the lower-noise single version, the NE5534, was temporarily deleted.

The same thing is going on everywhere; shareholder value is all that counts now. Quality be damned.

Supermarkets have been trying to shift consumers from cheap generic drugs to high-margin brand-name drugs for years. Why sell ibuprofen painkillers for 35p (the typical NE5532 price) when you can sell Nurofen Period Pain tablets for £4.69 (a typical expensive 'audio' op amp price)? The example is perhaps even worse in that both contain the same cheap active ingredient; at least the more expensive op amps are a different design.

I used to buy my NE5532s from Rapid, since they were the cheapest, but they told me that they are only selling the new TI version now.

Fake op amps

I've also come across a few fake NE5532s, normally provided with low-cost Chinese Hi-Fi project kits on eBay and the like. Why they think it's worthwhile faking a cheap (at the time) chip, I don't know. These are often renumbered 4558s; their low 3.5mA current consumption is the giveaway, as a real and proper NE5532 draws 8-9mA.

Another 'feature' is that the fakes don't have the back-to-back protection diodes across the two input pins, making them easier to damage. Interestingly, a YouTube video from Michael Beeny (**https://youtu.be/Jr-AXau7va4**) said that the new TI NE5532s omit these diodes as well. Fig.46 shows how to check it.

However, the new ones I bought from Rapid did have the diodes when I checked (I later confirmed they are the old version). Photo 40 shows some fake NE5532s from an AliExpress preamp kit. Oddly, they drew 2.6mA, but the current consumption took off rapidly at ±20V.

Also strangely, they had a 0.9V drop between the input pins, both ways on diode test. I know of no other op amp that has this. There's only one way to find out what it is: dissolve the encapsulation.

Overrated?

Some audio equipment makes use of the NE5532's high voltage rating; I certainly do in some of my designs, running them at ±21V for years with no failures to date. Using them at ±22V will give a headroom of +24dBu. One company (can anyone out there remember who?) even ran NE5532s at ±24V with an acceptable initial infant mortality rate.

Someone told me it was Hill Audio, who advertised their 1980s mixers as using 5532s throughout, but I checked and they ran them at a conservative ±17.5V. Rupert Neve's ISA110 mixer channel used ±15V, with NE5534s. Its high headroom of +26dBm was achieved by audio transformers, one being a specially made toroidal 'orange doughnut' step-up output transformer.

The wonderful Clive Green designed Cadac consoles employed in West End theatres definitely used ±20V and NE5532s throughout.

One has to be careful with the power supply when running close to the maximum. Overvoltage protection must be in place on the power supply in case a regulator goes short-circuit, letting through the higher unregulated voltage, or all the op amps will die. For this, I sometimes strap a 5W, 43V Zener diode across the power rails on the board.

The loss of one rail while the other is still running must also be prevented, since this can blow the protection diodes on the input pins, causing permanently high distortion. The old Kemo filter company trick of putting 560Ω resistors in series with the input pins works to protect them.

I recommend it for followers on the inverting pin and on the non-inverting pin for summing amps and inverters. These techniques are shown in Fig.47.

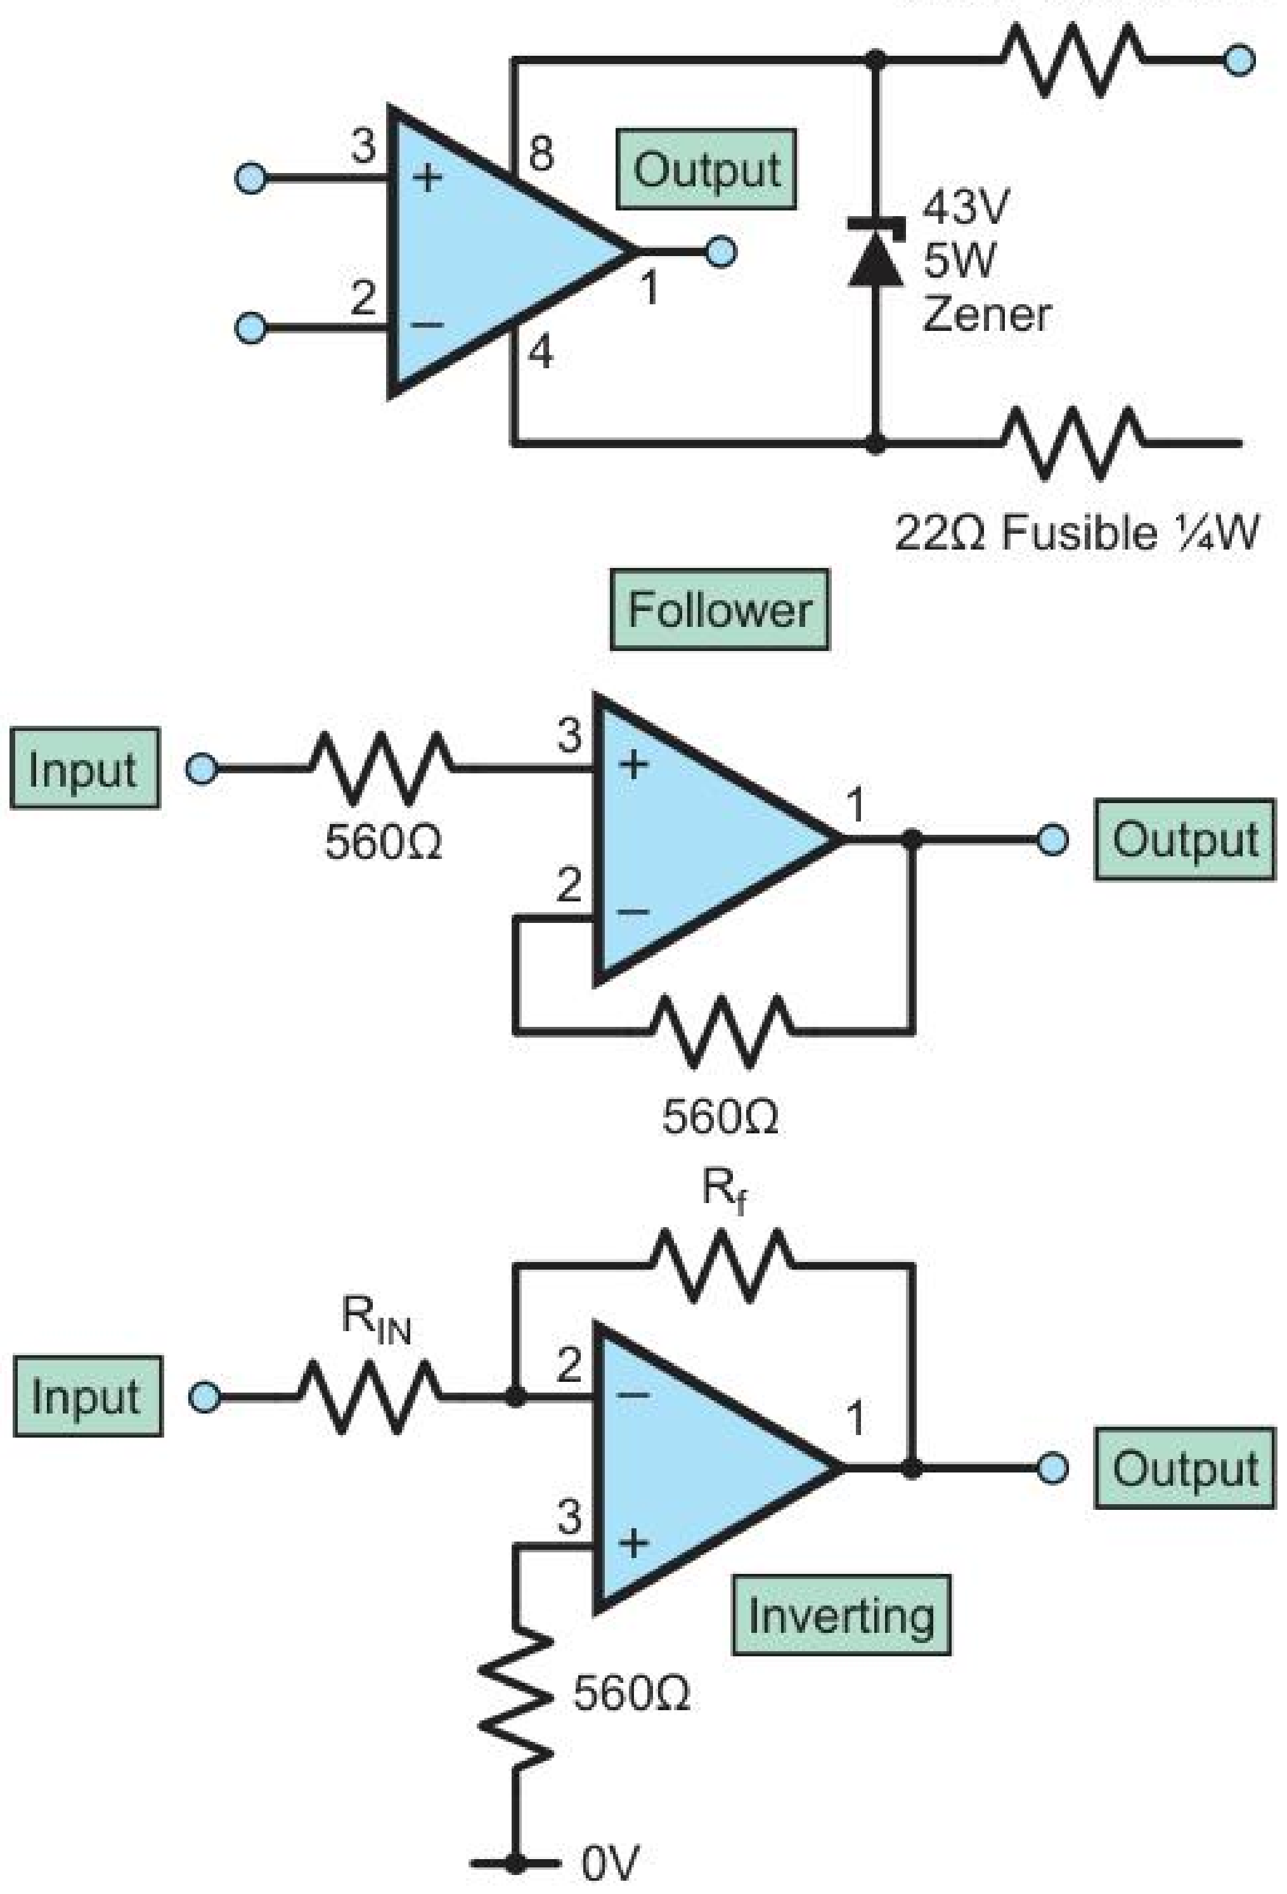

Fig.47: added protection for an op amp to prevent damage from overvoltage. The 560Ω resistors help if one power rail goes missing.

Back to Texas

Still, the idea of someone changing a chip for another with the same part number and it blowing up is totally irresponsible. The prospect of Les Misérables grinding to a halt due to a new Texas NE5532 being inserted as a replacement does not bear thinking about. (Sadly, the theatre desks are digital now because the old analog Cadac J series took the space of 10 seats selling at £200 a night).

The onsemi SMD versions of the NE5532 were downrated to ±20V because the dissipation was too much for the small package, and they stopped making them in 2022, although they are still available (for now).

It looks like TI's OPA2891 video op amp is a possible new NE5532 contender for the top audio amp, but it costs £5.00 and is over the top and so fast it will be unstable unless the circuit board is designed impeccably.

Craig informs me that there have been noise problems with some of the LM4562 (ex National Semiconductor) op amps, and they've gone up from £2.50 to £6.00 recently *[we've never managed to get them to work well; they result in much higher distortion when swapped into a circuit that's happy with an NE5532 or LM833, we suspect because they oscillate with less-than-perfect supply bypassing – Editor]*.

TI is even thinking of deleting these and replacing them with the LME49720, which looks identical, but of course costs £7.00. TI took over National Semiconductor's wonderful linear range, which I suspect is to be decimated.

National's LM833 (more on this later), which was always recommended as a slightly lower-cost NE5532 replacement where the high voltage rating was not required, also seems to have gone up in price.

At least the new TI NE5532s will still work at ±15V, which is the standard analog audio supply. Luckily, the noise spec is still the same at 5.5nV/√Hz. I will have to check the maximum output voltage swing into 600Ω, though, because it is now omitted from the data sheet. Rule 101 of data sheets: If a specification is not there, it's usually bad.

The Nisshinbo (formerly JRC) NL8802 at £1.35 with a ±22V rating looks like a viable NE5532 replacement contender, but it's only available in SMD packages. Nisshinbo is the only company still making the original device to the ±22V specification, including the 8-pin DIP device, the NJM5532DD, although it carries the worrying text, "not for new designs", on the data sheet.

Profusion Semiconductors, who are the UK distributor for Nisshinbo, only had the SMD version, the NJM5532CG-TE2 in stock recently, but they are getting a load of the 8-pin DIPs (NJM5532DD) in, before they cease production in March next year.

Profusion has been very helpful to me, giving me free THAT Corporation chip samples and real humans on the telephone, so I feel a bit of a plug is in order.

Back in phase, distortion

Looking at the distortion of the all-pass filter using different chips was interesting. The plot is shown in Fig.48. The 'new' NE5532s from Rapid were only slightly worse than the original old Texas devices.

They initially told me that I had purchased the new (2026) versions, but I later confirmed that they are still selling off stock of the pre-2026 version, which explains why these ones still perform well.

This is the problem with releasing a new part with the same model number as the old one!

I took a risk and tested a Rapid NE5532 (batch no 43A3K3M with no clear date code) at ±25V. There was no rapid rise in current, starting at 7.4mA at ±15V and holding at 8mA at ±25V. This pretty much confirmed they were the older version, especially, given what happened next.

Not long after that, I received a batch of TI NE5532s from Mouser. I've confirmed they have no diodes across the input pins and they blow up when a ±18V supply is connected. The current consumption is indeed lower, at 5mA with a ±15V supply.

The batch number is 59A7CLM and they have a little underlined G4 in the top right corner, which is a useful way of identifying these new parts. The peak distortion on the all-pass filter rises from 0.0007% to 0.0025% at 8V into 600Ω load using ±15V rails – a nearly four-fold increase in distortion from the 'same' part!

So the new NE5532s are definitely inferior and will blow up in some old circuits.

Strangely, the LM833 performed very badly in the all-pass filter at 8V RMS and 4V RMS, as shown in Fig.49. I think they must have been upset by the low impedance and high capacitance loading.

Another more subtle aspect of distortion is that generated by the frequency-determining capacitors: C6, C9, C15, C18, C21, C24 and C27. Polystyrene or polypropylene types give no problems. Using cheap polyester capacitors will give an increase, as shown in Fig.50. The distortion is 10 times worse at 8V RMS compared to polystyrene using 22nF 63V AVX 5mm polyester box capacitors.

[Editor's note: polystyrene, polypropylene and NP0/C0G ceramic capacitors can be considered 'ideal capacitors' for audio applications. Polyester types are definitely a step

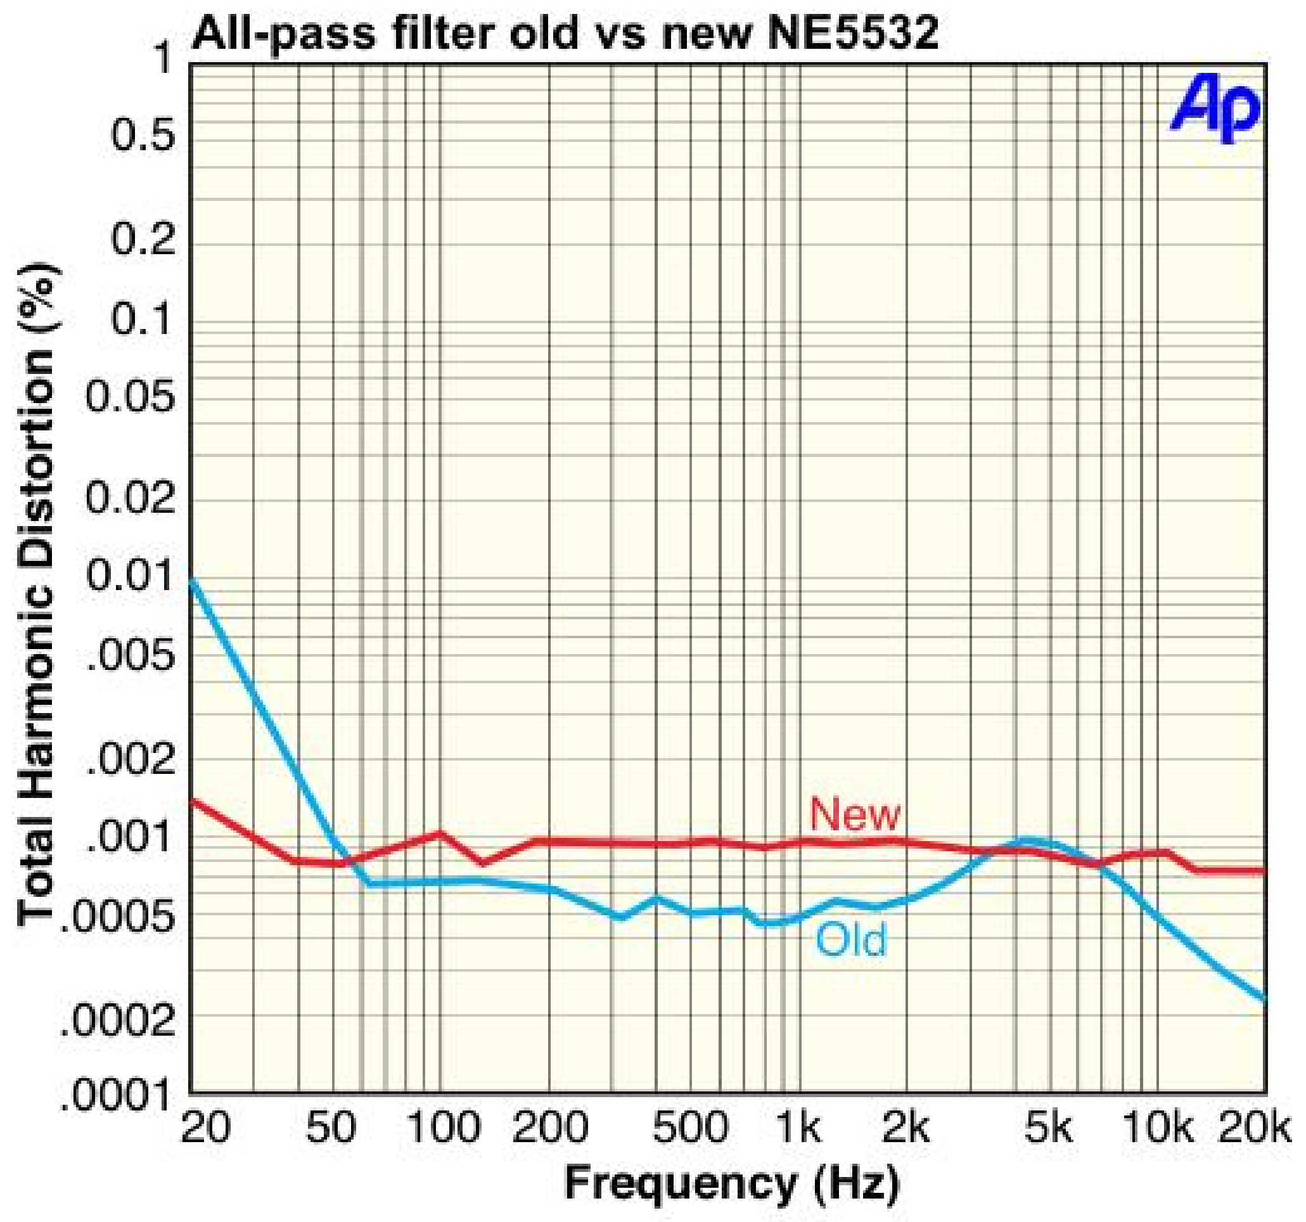

Fig.48: the new NE5532s from Rapid had slightly higher distortion but I don't think they are the latest version.

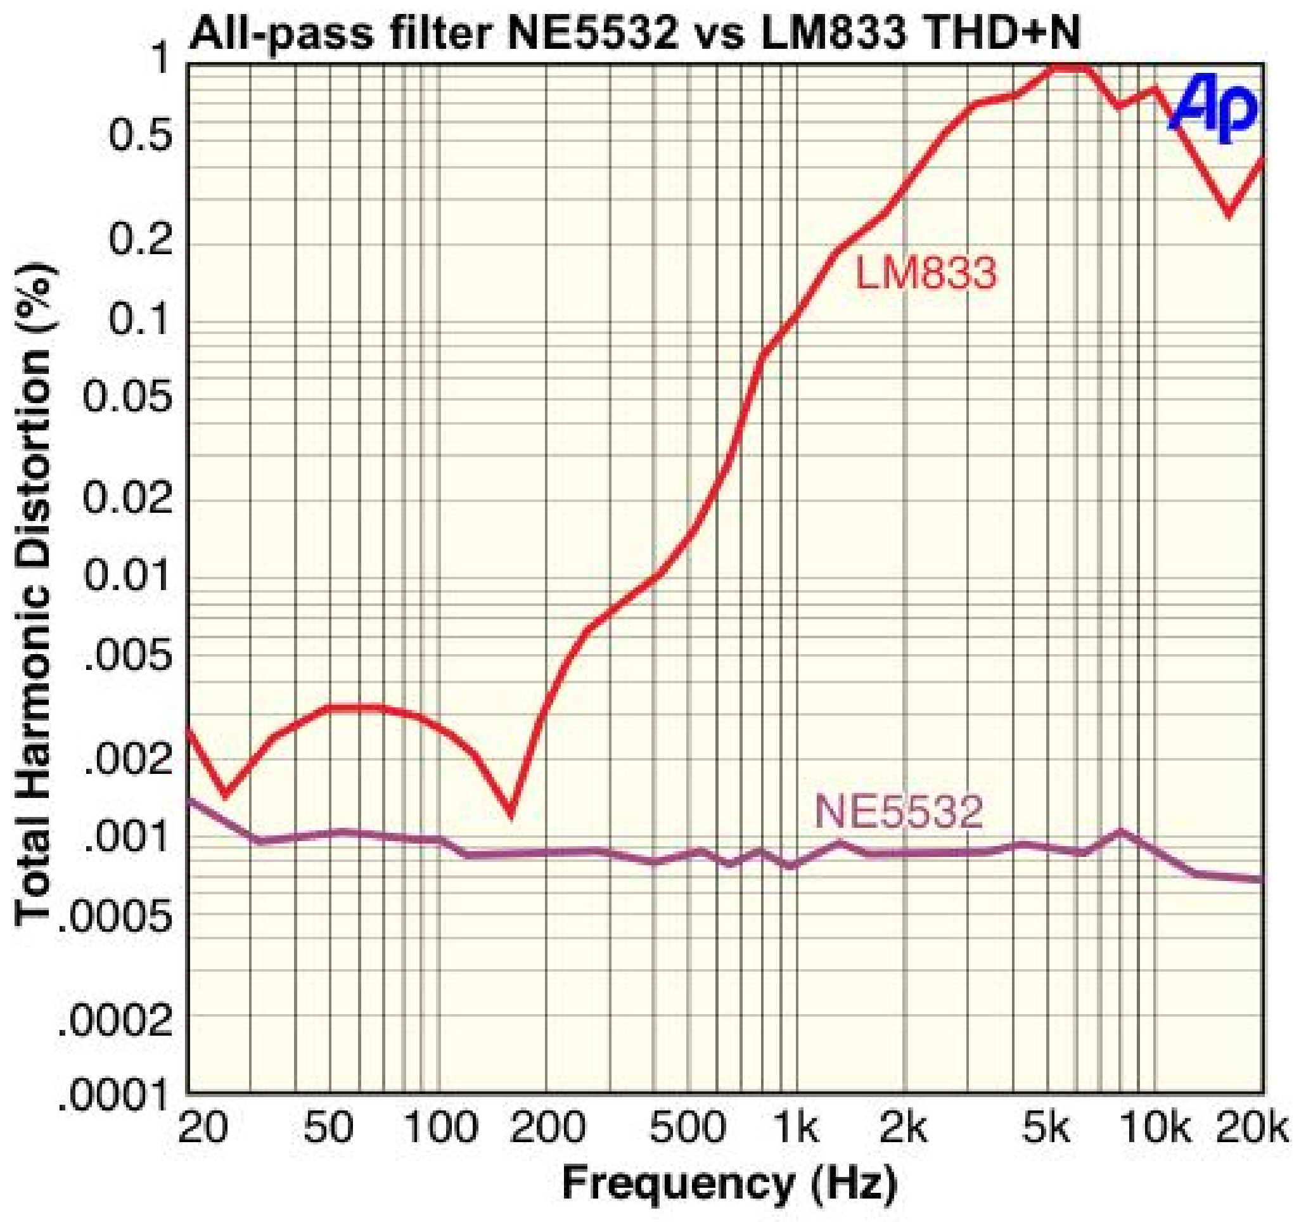

Fig.49: distortion with the LM833 is much worse than the NE5532 in this circuit, likely due to output loading.

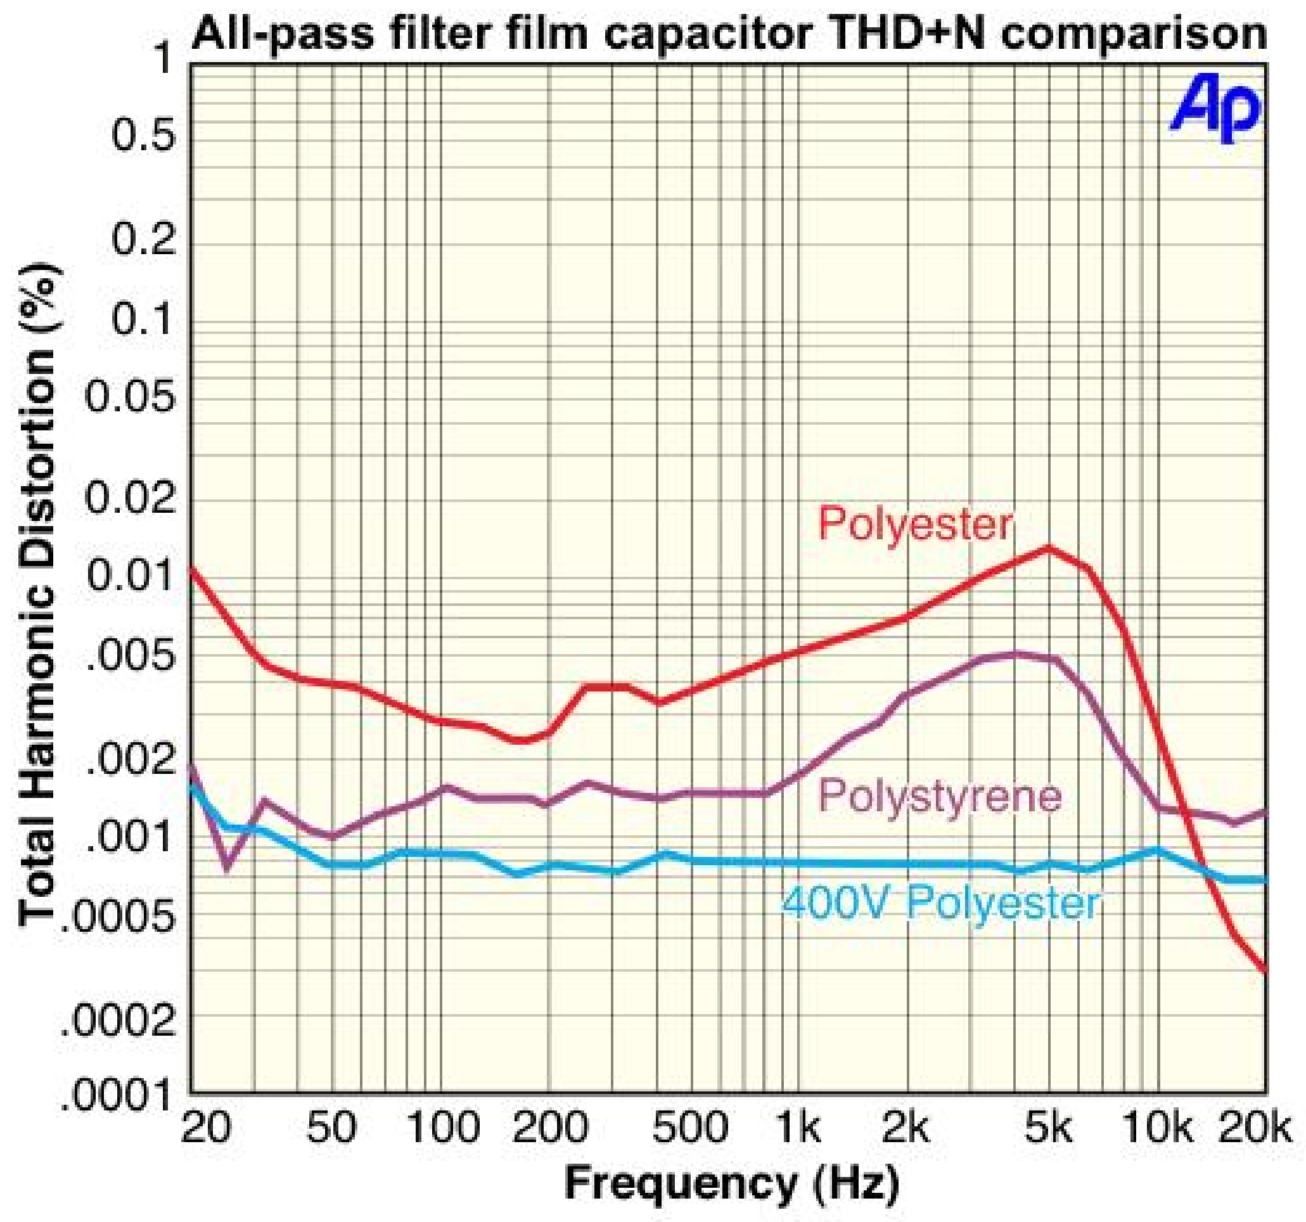

Fig.50: regular polyester filter caps are much worse than polystyrene but HV polyester caps are even better.

down; good enough for some tasks, but they are definitely not as linear, as demonstrated in Jacob's plot.]

With 22nF X7R ceramic capacitors, the distortion was almost 2%, even at 1V RMS, as shown in Fig.51.

[Editor's note: never use X7R, X5R, Y5V, Z5U or similar ceramic capacitors for audio; even for signal coupling. They have horrible non-linearity due to their high voltage coefficient. C0G/NP0 types are good for audio as they lack this fatal flaw. X5R/X7R types are good for supply bypassing, though.]

Clearly, this all-pass filter is a stressor of capacitors and op amps, as is to be expected. After all, it's a string of eight filters! I made a version using 400V 100nF Suntan dipped polyester capacitors (MPN TS02002G104KSB0D0R) from Rapid (order code 10-4167) and the distortion was fine (see Fig.50).

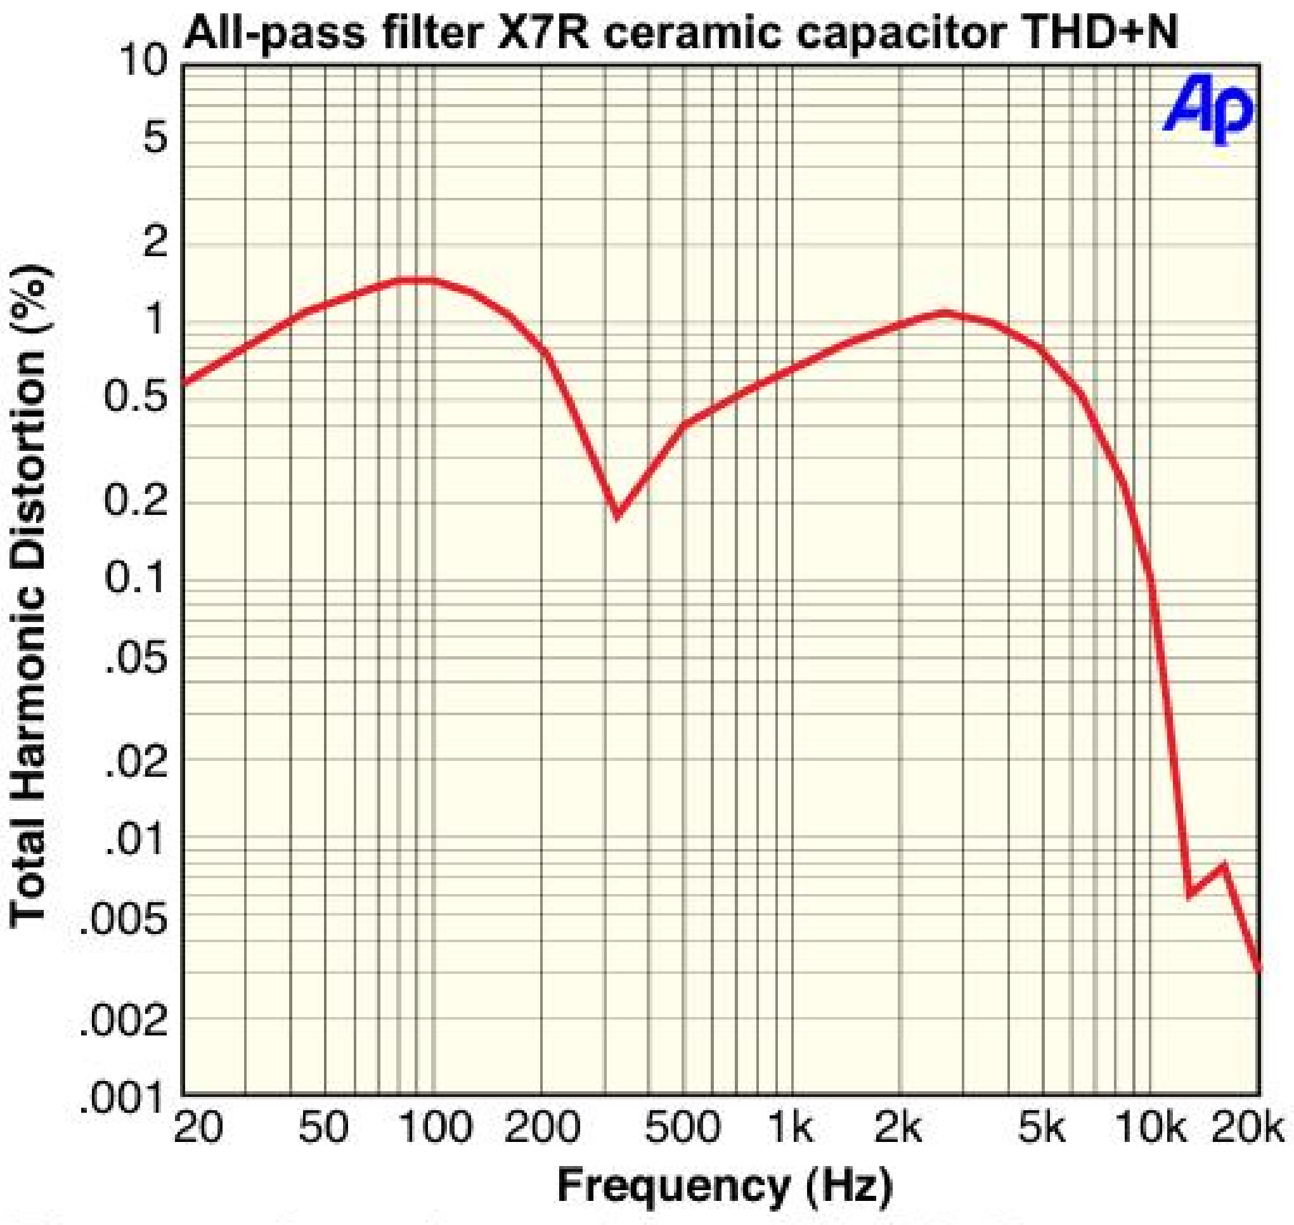

Fig.51: the distortion with X7R ceramic capacitors. What a mess! Only use them for supply bypassing.

The thicker 400V dielectric was suffering less stress. To safely mount these larger types of dipped capacitors, it's a good idea to provide strain relief on the leads, as shown in Photo 41. I hate to think how it would sound made with TL062 op amps and Y5V ceramic capacitors. Maybe it could be a new distortion effect; the Phaso-dilator, perhaps?

Losing the plot

With the horrific real UK inflation rate (only 3.5%? Who are they kidding?), it's time to do some hedging or 'prepping' of desirable components if you have spare cash. The huge price increases in solder and wire prices recently were a warning.

Low Value Import (LVI) duty relief below £135 on our Chinese component and PCB parcels to the UK is to be removed in March 2029. "Making Tax Digital (difficult)", with a possible software subscription, also worries me. I feel like falling on my soldering iron sometimes, but that would be a health and safety compliance failure!

Suntan®
2G104K

Photo 41: preforming capacitor leads using an RS tool. It's no longer available but Grove Sales PNR5000D (https://grovesales.co.uk/products/c-shape-forming-pliers) looks even better. I will try it out.

We'll continue the phase shuffle dance next month with putting boards in boxes and the phase-linear crossover board.

Thanks to my son, Harvey Rothman, for additional research in putting together this article. The PCB designs are by my usual partner in crime, Mike Grindle. **PE**

Teach-In 9 CD/DVD-ROM – Get Testing!

This series of articles provides a broad-based introduction to choosing and using a wide range of test gear, how to get the best out of each item and the pitfalls to avoid. It provides hints and tips on using, and – just as importantly – interpreting the results that you get. The series deals with familiar test gear as well as equipment designed for more specialised applications.

The articles have been designed to have the broadest possible appeal and are applicable to all branches of electronics. The series crosses the boundaries of analog and digital electronics with applications that span the full range of electronics – from a single-stage transistor amplifier to the most sophisticated microcontroller system. There really is something for everyone!

Each part includes a simple but useful practical test gear project that will build into a handy gadget that will either extend the features, ranges and usability of an existing item of test equipment or that will serve as a stand-alone instrument. We've kept the cost of these projects as low as possible, and most of them can be built for less than £10.

FREE CD-ROM
ELECTRONICS TEACH-IN 9
£8.99
FROM THE PUBLISHERS OF EVERYDAY PRACTICAL ELECTRONICS
GET TESTING!
Electronic test equipment and measuring techniques, plus eight projects to build
FREE CD-ROM TWO TEACH-INs FOR THE PRICE OF ONE
Multimeters and a multimeter checker
Oscilloscopes plus a scope calibrator
AC Millivoltmeters with a range extender
Digital measurements plus a logic probe
Frequency measurements and a signal generator
Component measurements plus a semiconductor junction tester
PIC n' Mix
Including Practical Digital Signal Processing
PLUS...
YOUR GUIDE TO THE BBC MICROBIT
A LOW-COST ARM-BASED SINGLE-BOARD COMPUTER
Teach-In 9
Get Testing

The CD/DVD-ROM includes the article series in PDF format, plus the software for the *PIC n' Mix* series of articles. It also includes the full *Teach-In 2* book – *Using PIC Microcontrollers – A practical introduction* – in PDF format.

JUST £8.99

(includes P&P to UK if ordered direct from us)

ORDER YOUR COPY TODAY: www.electronpublishing.com

AOShop

The home for specialist audio, analogue and historic components – provided by Jake Rothman, *PE’s Audio Out* columnist. The AOShop is your best bet for classic analogue ‘NOS’ (new old stock) components, including all parts for *Audio Out* projects and designs.

Ge semiconductors

Small-signal PNP transistors	
NKT214F, OC57, OC59, 2N1377, 2N525	£1.00
Low-noise PNP transistors	
GET106	£2.50
Small power PNP transistors	
AC153, AC153K, AC188	£1.50
Small power NPN transistors	
AC176, AC141K	£2.00
High-voltage PNP transistors	
OC77, CV7001	£1.50
TO3 PNP power transistors	
OC22, CV7054 (OC23), OC25, OC35, OC36, AD143, AD149, AD161, AD162	£2.00
AD140	£3.50
AD149, AD161/2 matched pair	£5.00
RF PNP transistors	
OC41, OC42, 2SA12, 2SA53,AF124, AF178, GET872A	£1.50
NPN transistors	
OC139, OC140, ASY73	£2.00
Diodes	
CV7049 (OA10), CG92 (OA91)	£0.50

Si semiconductors

Diodes	
ZC5800 RF Schottky	£0.20
Low-noise silicon transistors	
BFW16A, 2SC3071, 2SC3068, 2SA1016K, 2SC2362K, 2SA970BL	£1.50
2SC2204, 2SD655, BC550C	£0.50
ZTX651	£0.30
ZTX751	£0.50
RF transistors (suitable for Theremin)	
BF199	£0.50
Audio power MOSFETs	
Exicon 10N20, 10P20	£6.50
Hitachi 2SJ99, 2SK343	£3.50
Hitachi 2SJ56, 2SK176	£8.50
JFETs	
BFW11, BFW10, TIS73L, J177, J113, U1994, U1898, 2SJ176, J201	£1.00
J175, J176, J112, J111, 2N3820, 2N5457, BF244, 2N5460, J230	£0.60
Small power output/driver transistors	
2SB649A, D669A, 2SA1208, 2SC2910, MJE253G, MJE243G, 2SA1725, 2SC4511	£1.20
BD139,BD140, BD135, BD136, BD435, BD436, MPSA63, BCV46	£0.50
MOSFETs	
ZVP2106A	£0.30
Dual-gate 3SK45, BFS28	£1.50
Metal-cased transistors	
BCY71	£0.30
Amplifiers	
LM318 high-speed op amp	£0.35
µA709 metal-cased op amp	£2.00
5534H metal-cased op amp	£3.00
LM384, TDA2030A, TDA2050V	£1.50
TAA435 (Mullard power amp driver for AD161/AD162)	£1.20

ICs

That/dbx 2180 VCA/VCF	£6.00
CA3080 VCA/VCF (vocoder)	£2.80
CA3280 VCA/VCF	£4.00
LM13600/ LM13700 VCA/VCF	£2.00
CA3086 transistor array	£1.00
PT2399 echo/delay	£1.50
Original TI 44V NE5532	75p
Nisshimbo 44V NE5532DD	£1.00
SMT version	50p
New 36V TI NE5532	50p
Original TI NE5534A	£1.40

Dual transistors

2N2639, 2N2223, 2N2910 (NPN)	£4.00
E401 (JFET Moog)	£4.00
2SK2145-Y dual JFET	£0.80
2N5564 JFET	£8.00
HN3C51F, HN3A51F	£1.00
DMMT3904/6, HN1A01F, HN1C01F	£0.50

Loudspeakers

PE Mini-Monitor	
Volt PE165 6.5-inch woofer (each)	£85
Morel MDT29 tweeter (each)	£25
Kit *pair* of PE165/MDT29 plus Volt crossover parts and PCBs	£299
Monacor DT-28N tweeter (each)	£35
Vifa 19mm BC20SC15-04 tweeter (each)	£15
Volt crossover inductors	
1.2mH, 1.5mH, 2mH, 2.7mH, 0.5mH (tapped at 0.3mH) (each)	£5.85
Fully tested reclaimed speakers	
Vifa BC14 5-inch woofer (each)	£10
Vifa TC26 1-inch tweeter (each)	£10
Low-price speakers	
Philips 4-inch 4070	£2.00
EMI 10x6-inch, 30Ω Alnico	£7
64mm 64Ω neodymium	£1.20

Capacitors

Note ‘10/63’ denotes ‘10µF 63V’.

Polyester	
3.3/100, 4.7/250, 4.7/63	£1.00
5.6/63, 8.2/63, 10/63	£2.00
Mullard ‘Mustard’ C296 0.22/400	£2.00
Polycarbonate	
Axial 2.2/63 1%, 4.7/160, 6.8/63	£1.00
Radial	
6.8/160V, 10/63	£2.00
Reclaimed 22/63	£2.00
Polystyrene	
Philips 1% 4.7nF/160, 6.2nF/500, 12nF/63, 22nF/63, 110nF/63, 24nF, 2nF	£1.00
RIFA 2% 100nF for active xov.	£1.00
Suflex 90.9nF 0.5%	£2.00
Suflex 2.5% 10nF/63 (rad. or ax.)	£0.50
Radiation resistant	
Siemens cellulose acetate MKL 2.2/25	£0.80
Electrolytic – Mullard blue 017 series	
10/25, 22/25, 100/10	£0.50
150/40, 470/40, 1000/40	£1.00
Tantalum – axial metal cased	
22/50, 47/35, 68/25, 100/20, 120/10, 150/16, 220/10, 330/6	£2.00
22/35, 33/35, 47/20, 68/15, 100/10, 150/6	£1.25
4.7/50, 6.8/35, 10/25, 10/35, 22/15	£1.00
Axial moulded-case tantalum	
Kemet axial 6.8/10	£0.30
Kemet radial 33/10	£0.40
STC radial 100/20	£1.50
Tantalum bead	
22/50	£1.00
470/3	£2.00
680/6.3	£3.00
Wet tantalum	
220/25 axial	£3.20
Castanet button 140/30, 470/3	£3.20
Hughes 540/10	£3.20
Bipolar	
Hermetic bipolar tantalum 16/35	£3.20
Elcap axial 10/50	£0.50
Generic radial 100/16, 470/35, 100/ 63, 22/35, 4.7/35, 220/16	£0.50
Philips solid-aluminium (axial) 121/123	
47/16	£1.00
330/6.3	£2.00
100/35	£4.00
Philips Pearl 122 series (radial)	
10/16	£0.30
Silvered mica (radial)	
1nF/500 1%	£1.00
Trimmer capacitors	
Vishay plastic-film 4-40pF	£0.80
Vishay plastic-film 5.5-45pF	£1.00
Vishay plastic-film 5-80pF	£1.20
Vishay plastic-film 10-250pF	£2.00
Mica 1-12pF, 2-40pF	£1.00

Audio transformers and inductors

Eagle transformers	
LT44, LT722 driver, LT700, LT723 500Ω output	£2.50
LT30 500mW output	£3.50
5:1 interstage (fuzz box)	£1.50
Repanco	
T/T3 splitter transformer	£4.00
CH2 5mH RFC	£2.00
Balanced output transformer	
Vigortronix 600Ω VTX-101-007	£10
Vigortronix 600Ω VTX-101-3001	£10
Vigortronix 600Ω VTX-101-3002	£15
Gardners 150Ω	£10
Reclaimed BBC LL74/MPC nickel core 600Ω	£12
Low-level 1:1 600Ω (modem)	£2
Reclaimed mic input transformer	£15
Inductors	
82µH, 4.7mH, 100µH, 270µH, 10µH, 14µH (low Z)	£0.50
4-inch ferrite rod with MW and LW windings	£3.00

Special resistors

Bourns wire-wound trimmer 5kΩ 3059 JM panel-mount	£2.00
Thermistor RA53, R13	£4.00
A13	£2.00
Thermistor CZ1, CZ6	£1.50
Holco H2 2.2MΩ 1W, 1%	£1.00
Welwyn 1GΩ 2W	£1.00
5k Bourns 3321H cermet trimmer	£0.50

Potentiometers

Bourns 81	
25kΩ lin cermet	£2.00
25kΩ lin, 5kΩ lin cond. plastic	£3.00
Bourns 91	
10k dual-gang lin or log	£10.00
Plessey moulded-track	
5kΩ log with switch	£3.00
50kΩ A/log	£3.00
Bourns 8-gang 50k A/log	£7.50
ALPs 4-gang 100k A/log	£3.00
Mil 250kΩ lin dual	£5.00
Alpha	
16mm 4.7kΩ A/log	£0.80
220kΩ A/log	£0.80
10kΩ lin centre-detent dual-gang	£1.50
Allen Bradley J series/Honeywell	
10kΩ lin 1 million cycles	£5.00
Blore Edwards AB 45 dual 5kΩ A/log with switch	£3.50
Alps RK9 dual-gang 5kΩ RD law	£4.00
Alps 50k log tapped motor. stereo	£5.00
BI P260 500kΩ log conductive plastic 1 million cycles	£2.00
Colvern wire-wound 100kΩ or 50kΩ dual-gang 3W	£5.00
Mil Spec hermetic 10Ω	£8.00

Miscellaneous

PCBs	
8-gang All-pass Filter PCB	£3.00
8-gang Pot PCB	£1.50
Polarity Inversion PCB (Alps)	£1.50
Polarity Inversion PCB (alt.)	£1.50
All-pass Filter PCBs (three)	£4.00
.. Rotary Switch (AO0626ROT)	£1.50
.. Pot Board (AO0626POT)	£1.50
Silent Stomp Switch (pair)	£4.00
Penta-fuzz Bender (*PE*, Feb26)	£3.00
Pre-driver (*PE*, Jan26)	£3.00
Pre-driver jack bd (*PE*, Jan26)	£1.50
Fuzz Box (*PE*, Nov25)	£3.00
Fuzz Box jack bd (*PE*, Dec25)	£1.50
Dual JFET Buffer (*PE,* Oct25)	£2.00
Mosfet Buffer (*PE,* Oct 2025)	£2.00
Dual Mosfet Buffer (*PE,* Oct25)	£2.00
Cascode Amp/PS (Sep 2025)	£2.00
Pocket Theremin (*EPE*, 1996)	£2.00
Elysian Theremin (*EPE*, 1996)	£6.00
Synth VCF, VCO (*EPE*, 2017)	£3.00
48V PSU (*EPE*, 2019)	£3.00
SMT dual transistor adapter	£0.40
Switches Blore Edwards/Plessey/etc	
ALPs 4-pole push switch	£1.50
MH1 4-pole 3-way	£3.00
Elma 01 2-pole 6-way	£6.00
3PDT guitar pedal stomp switch	£3.50
Valves	
EC88, PCL84	£3.00

Contact

Jake Rothman
Manchester House
Market Street, Craven Arms
Shropshire SY7 9NN
(visit by appointment)

+44 (0)1597 829102

jrothman1962@gmail.com

Payment

PayPal, bank transfer, cheques (payable to “J Rothman”, UK pounds only). No VAT payable

Postage

Small Jiffy bag £2.99
Small package £4.99

USB Solar Charging System

Simple Electronic Projects with Julian Edgar

Charge and power all your USB devices from solar with this low-cost system. It's inexpensive to put together, and once you've built it, charging your devices won't cost a cent!

I have lots of solar-powered devices. A solar-powered smart watch, a solar-powered iPhone, solar-powered noise-reducing headphones and a solar-powered mini floodlight. In fact, every one of my devices charged from a USB adaptor is now solar powered. This is achieved very simply and, if you're careful with your purchasing, very inexpensively too.

The parts required

Only a few parts are required: a solar panel (Photo 1), a charge controller (Photo 5), a 12V battery and a one or more 12V to 5V USB converters (Photos 2, 3 & 4). You could spend hundreds of dollars on assembling these parts – or you could do as I did, and buy mostly second-hand, via online market places and/or make use of parts that others have thrown away.

The solar panel needs to have an output voltage suitable for charging a 12V battery through a charge controller. This means getting a panel that has a maximum open circuit voltage of about 18-30V, depending on the selected controller.

The maximum solar panel output power will also be governed largely by the controller you select. Small charge controllers may have a 60W limit, while larger controllers are good for 250W. So selecting the correct panel is done in conjunction with the charge controller you've picked.

Second-hand solar panels are now ridiculously cheap – expect to pay from about £12 for a suitable used panel. Also remember that, in this application, the panel's original maximum output power is probably not needed. To put that another way, this is a good project for reusing degraded panels that otherwise would go to scrap.

The appropriate charge controller will also depend on the battery type you're going to use. Sealed lead acid (SLA) batteries are ideal for this application as they can be mounted inside without concern for acid spills or venting of gases. However, SLA batteries tend to be expensive – even second-hand – so you may wish to use a conventional car battery.

A major benefit of using a car battery is that you can get one free of charge. Simply visit a local mechanic or car battery supply shop. There you will find literally dozens of batteries that have been discarded – they've been replaced as no longer being suitable for cranking engines.

Starting an engine is very demanding on a battery; the current draw is in hundreds of amps. So these batteries often still have sufficient capacity to work as a storage battery in a solar system of the type being covered here.

Photo 1: I mounted this 100W solar panel on a disused satellite dish on the roof. The panel can also be wall-mounted, or even just anchored to the ground. Second-hand solar panels are now very cheap.

5V
2.1A
12.5
5V
2.1A
Photo 2: this 12V-to-USB converter/charger provides two outputs and a voltmeter to allowing monitoring of system voltage. It cost about $11 from Banggood (similar units can be found on AliExpress & eBay).

Solar panel and battery ratings

The greater the power of the solar panel, the better it will work through poor weather and the more battery charge you will have to work with. Also, the larger the battery, the longer the system will cope with cloudy days when little solar output is available.

To give you a guide, where I live about 100km north of Canberra, Australia (at a similar latitude to Gibraltar), in four years I have never gone close to running out of power using a 100W panel and a 26Ah SLA battery – the latter bought as defective and so probably having only half this nominal capacity. That includes charging power tool batteries as well as my phone, watch, camera, etc.

Note also that smaller panels and batteries tend to be more expensive second-hand, so there's a further advantage in going big.

Photo 3: this non-switched panel with two USB outputs costs about $6 from AliExpress. Ensure you get a device that runs from 12V.

When selecting a battery from the discard pile, use a multimeter to find a battery that still has an open-circuit voltage above 12V. If you are going to use the battery inside, select one that is fully sealed.

A battery with flat terminals to which lugs can be bolted will be easier to wire than a battery with round terminal posts. If you have one with round posts, you'll need to get matching terminals to attach wires to it, which will be an extra cost.

Once you have selected a solar panel and battery, you can pick a charge controller to suit. Here it's worthwhile buying new. At the time of writing, a 30A (about 400W) 12V charge controller costs around £6, including freight. They can be bought from AliExpress and similar suppliers. Ensure the controller can be configured to suit the battery type you're using – most can.

Many controllers also have built-in USB 5V outputs. If you choose to use these, you don't need to buy an additional USB converter. However, more for convenience and appearance, I added three USB output converters to my system. Each of these has two USB outputs (giving six in total), on/off switches and an onboard LED display showing the battery voltage. These units are about £5.50 each.

Note that when selecting these, ensure you don't get ones designed to plug into a cigarette lighter socket – you want wired-in ones.

If you want an even cheaper approach, just buy a 12V-to-USB wired converter that comprises a sealed box, USB output socket on a lead, and input power connections.

In addition to these parts, you will need assorted cabling, terminals, an inline fuse holder and fuse, and possibly a panel on which to mount the USB outlets.

Building the system

Before starting to build the system, consider the cabling requirements. I have the charging system working in my home office, so I had to run a cable from the roof-mounted solar panel through the wall of my house. In my case, that was easy but, in many houses, that will be quite hard! If that's the situation, consider having the charging system in an outside workshop or shed.

It will still be very useful there – for example, most battery-powered tools have 12V car chargers available for their batteries, so the system can be used to charge power tool batteries.

Also, consider how the solar panel will be mounted. In my situation, a disused satellite TV dish antenna was on the roof, pointing towards

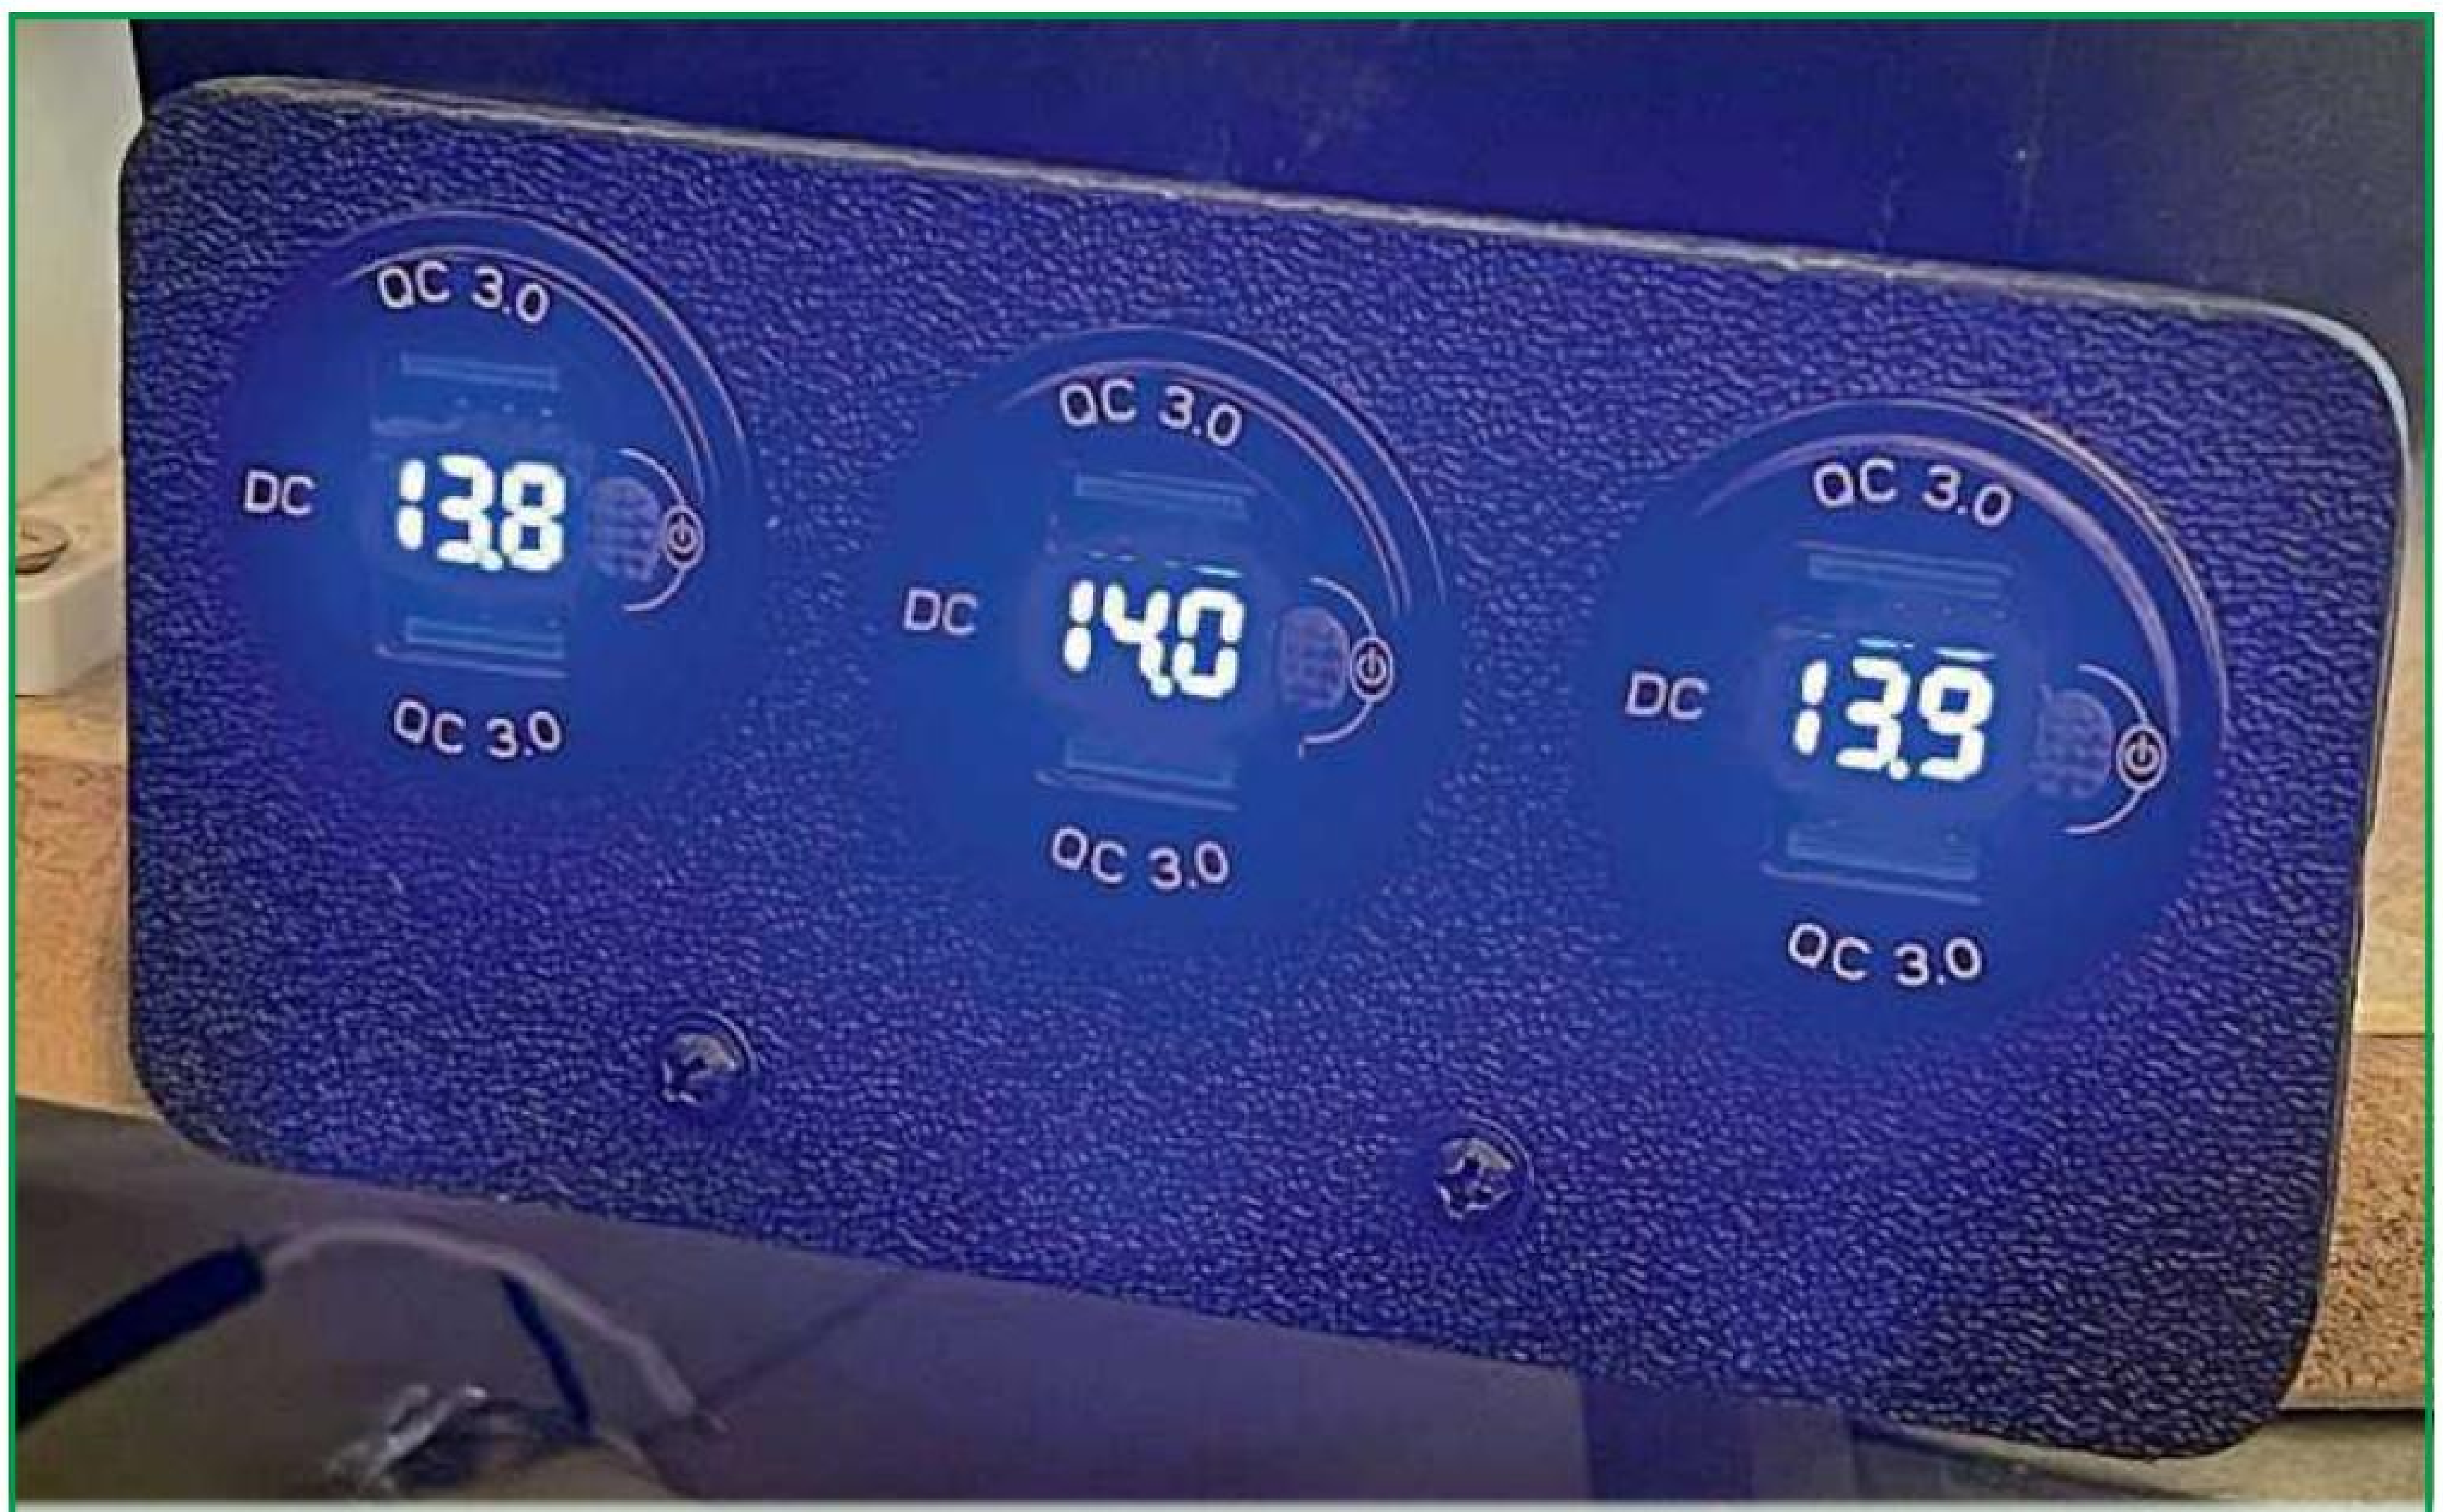

Photo 4: in my system, three 12V-to-USB adaptors provide six outlets at the desk in my home office. Note the inconsistency in the voltage readouts – at these prices, you can't expect perfection! This is the charging voltage on a sunny day.

SOLAR CHARGE CONTROLLER
88.8 V
USB

Photo 5: this solar panel charge controller from AliExpress incorporates two USB outputs so, if you wish, you can charge your items directly from this module.

the equator. Mounting the panel was just a case of attaching the panel to the dish – quick and easy! However, again, that may not be the situation in your case. If roof-mounting is difficult, consider mounting the panel on a wall or even on the ground.

The panel should face towards the equator and be tilted at an angle that approximately corresponds to your latitude, although horizontal panels will generally work OK. We don't need to squeeze every drop of power out if it!

Fig.1 shows the wiring – it is very simple. Ensure you place the fuse close to the battery; it should be rated at the maximum charging current, as dictated by the controller.

If you are not used to working with 12V storage batteries, keep in mind that although the voltage is low (so you won't get a shock), the ability to deliver current is very high and so you must be careful to ensure that the battery is never short-circuited. I have seen this done when someone inadvertently dropped a spanner across the battery terminals... not good!

Always take great care when attaching battery connectors; the battery terminals must be insulated when the battery is in use. An easy way to achieve this is to place the battery in a dedicated box. Boxes designed to house car batteries are available from about £7.50 from local suppliers.

Parts List – USB Solar Charging System

1 12V solar panel
1 12V charge controller (type to suit panel and battery)
1 12V rechargeable battery
1 pair of battery terminals (may not be required depending on battery type)
1 inline fuse, rated to suit charge controller
1 or more 12V-powered USB chargers
various lengths of wire, rated to handle the maximum charging current

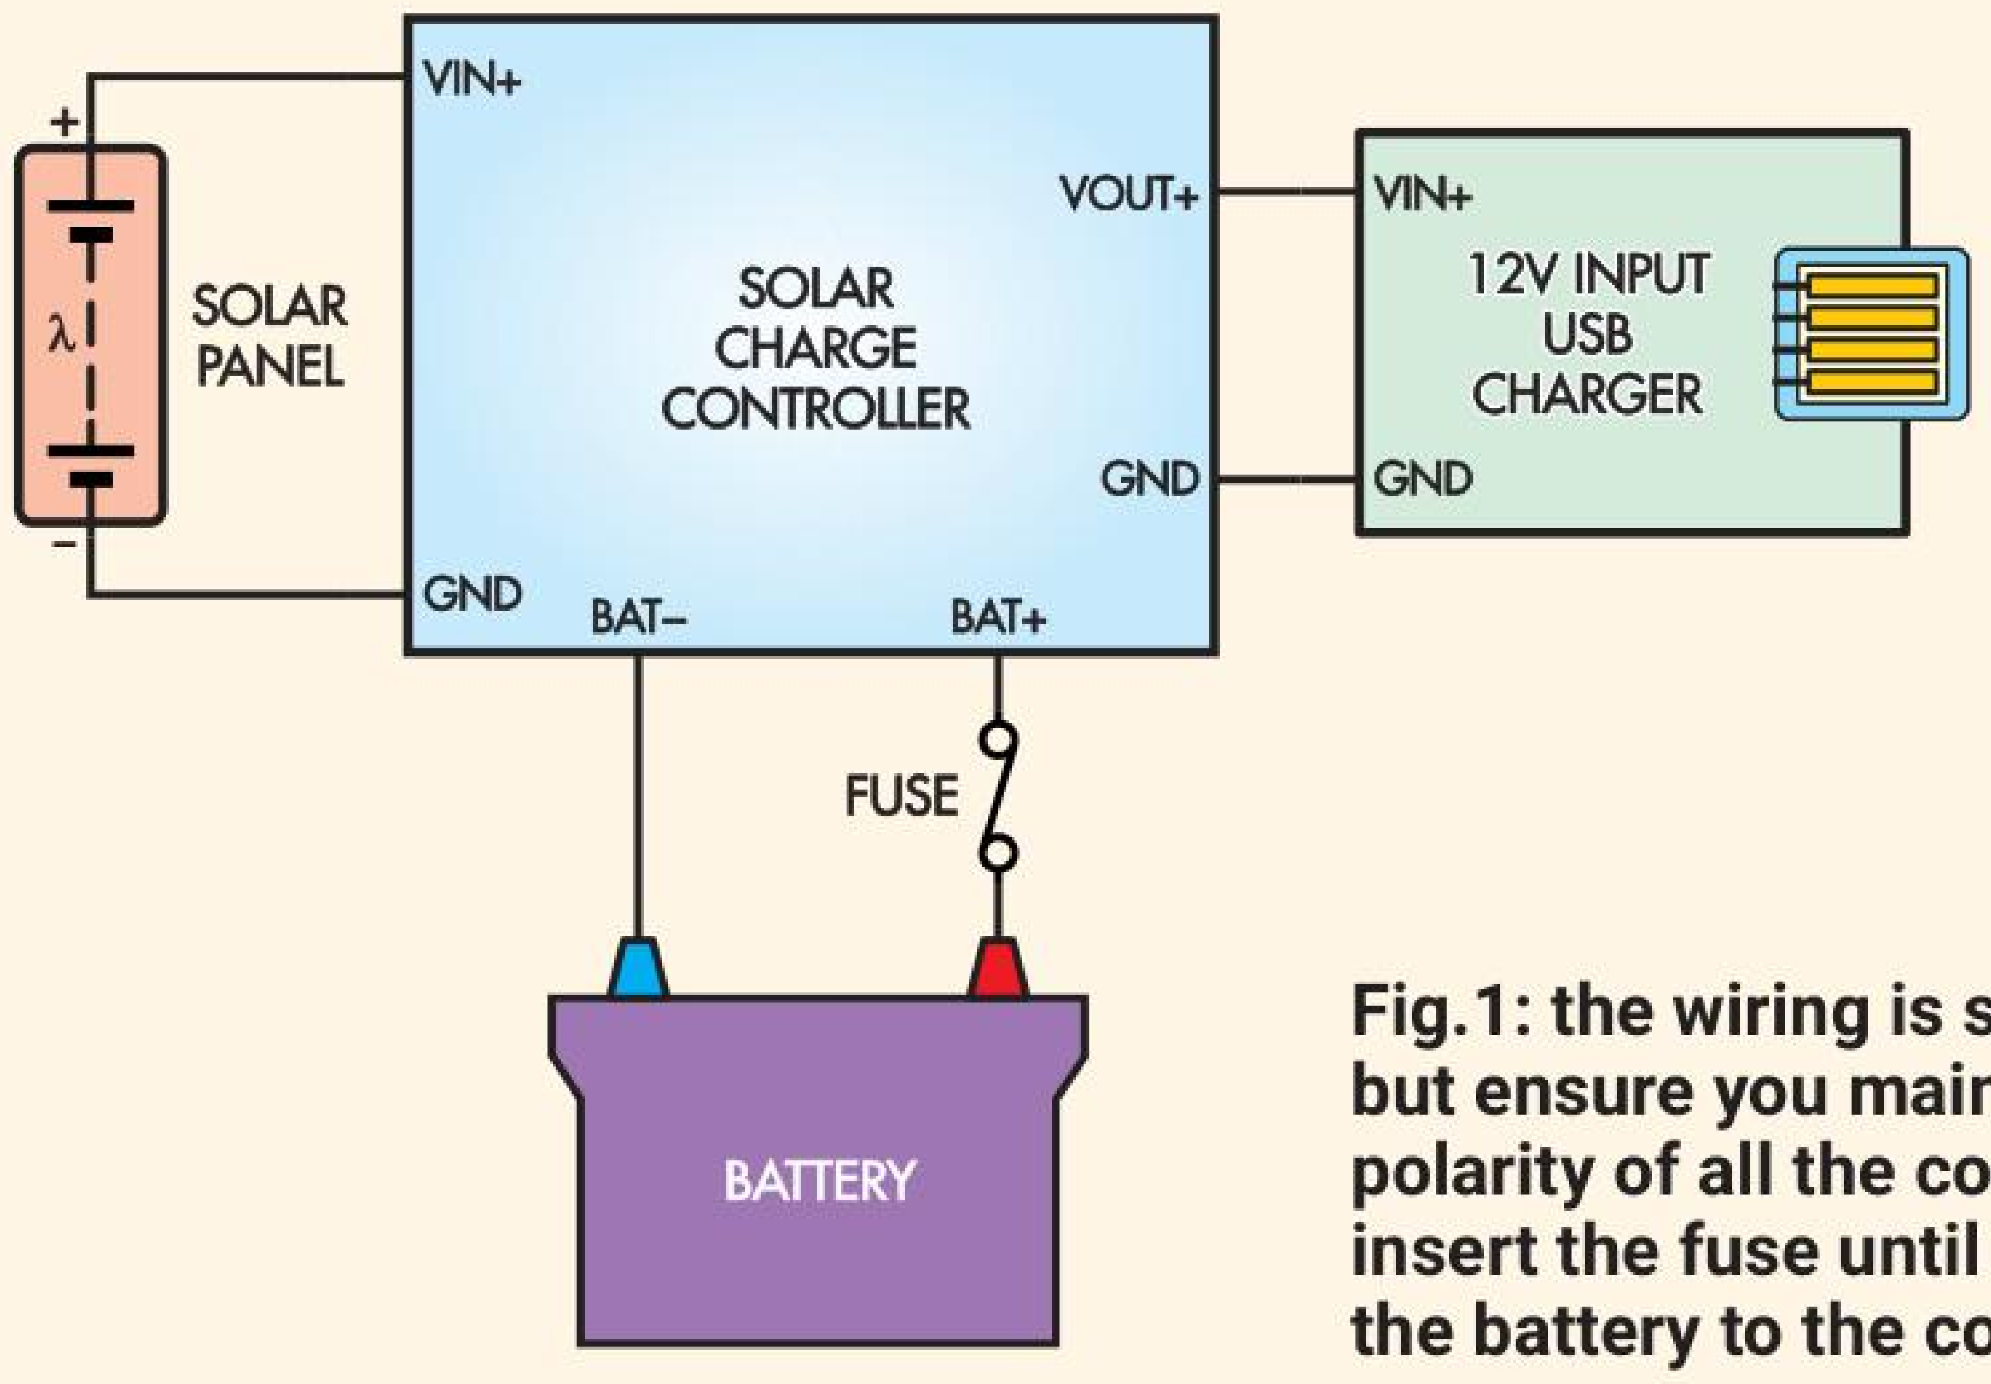

Fig.1: the wiring is straightforward, but ensure you maintain the correct polarity of all the connections. Don't insert the fuse until you have wired the battery to the controller.

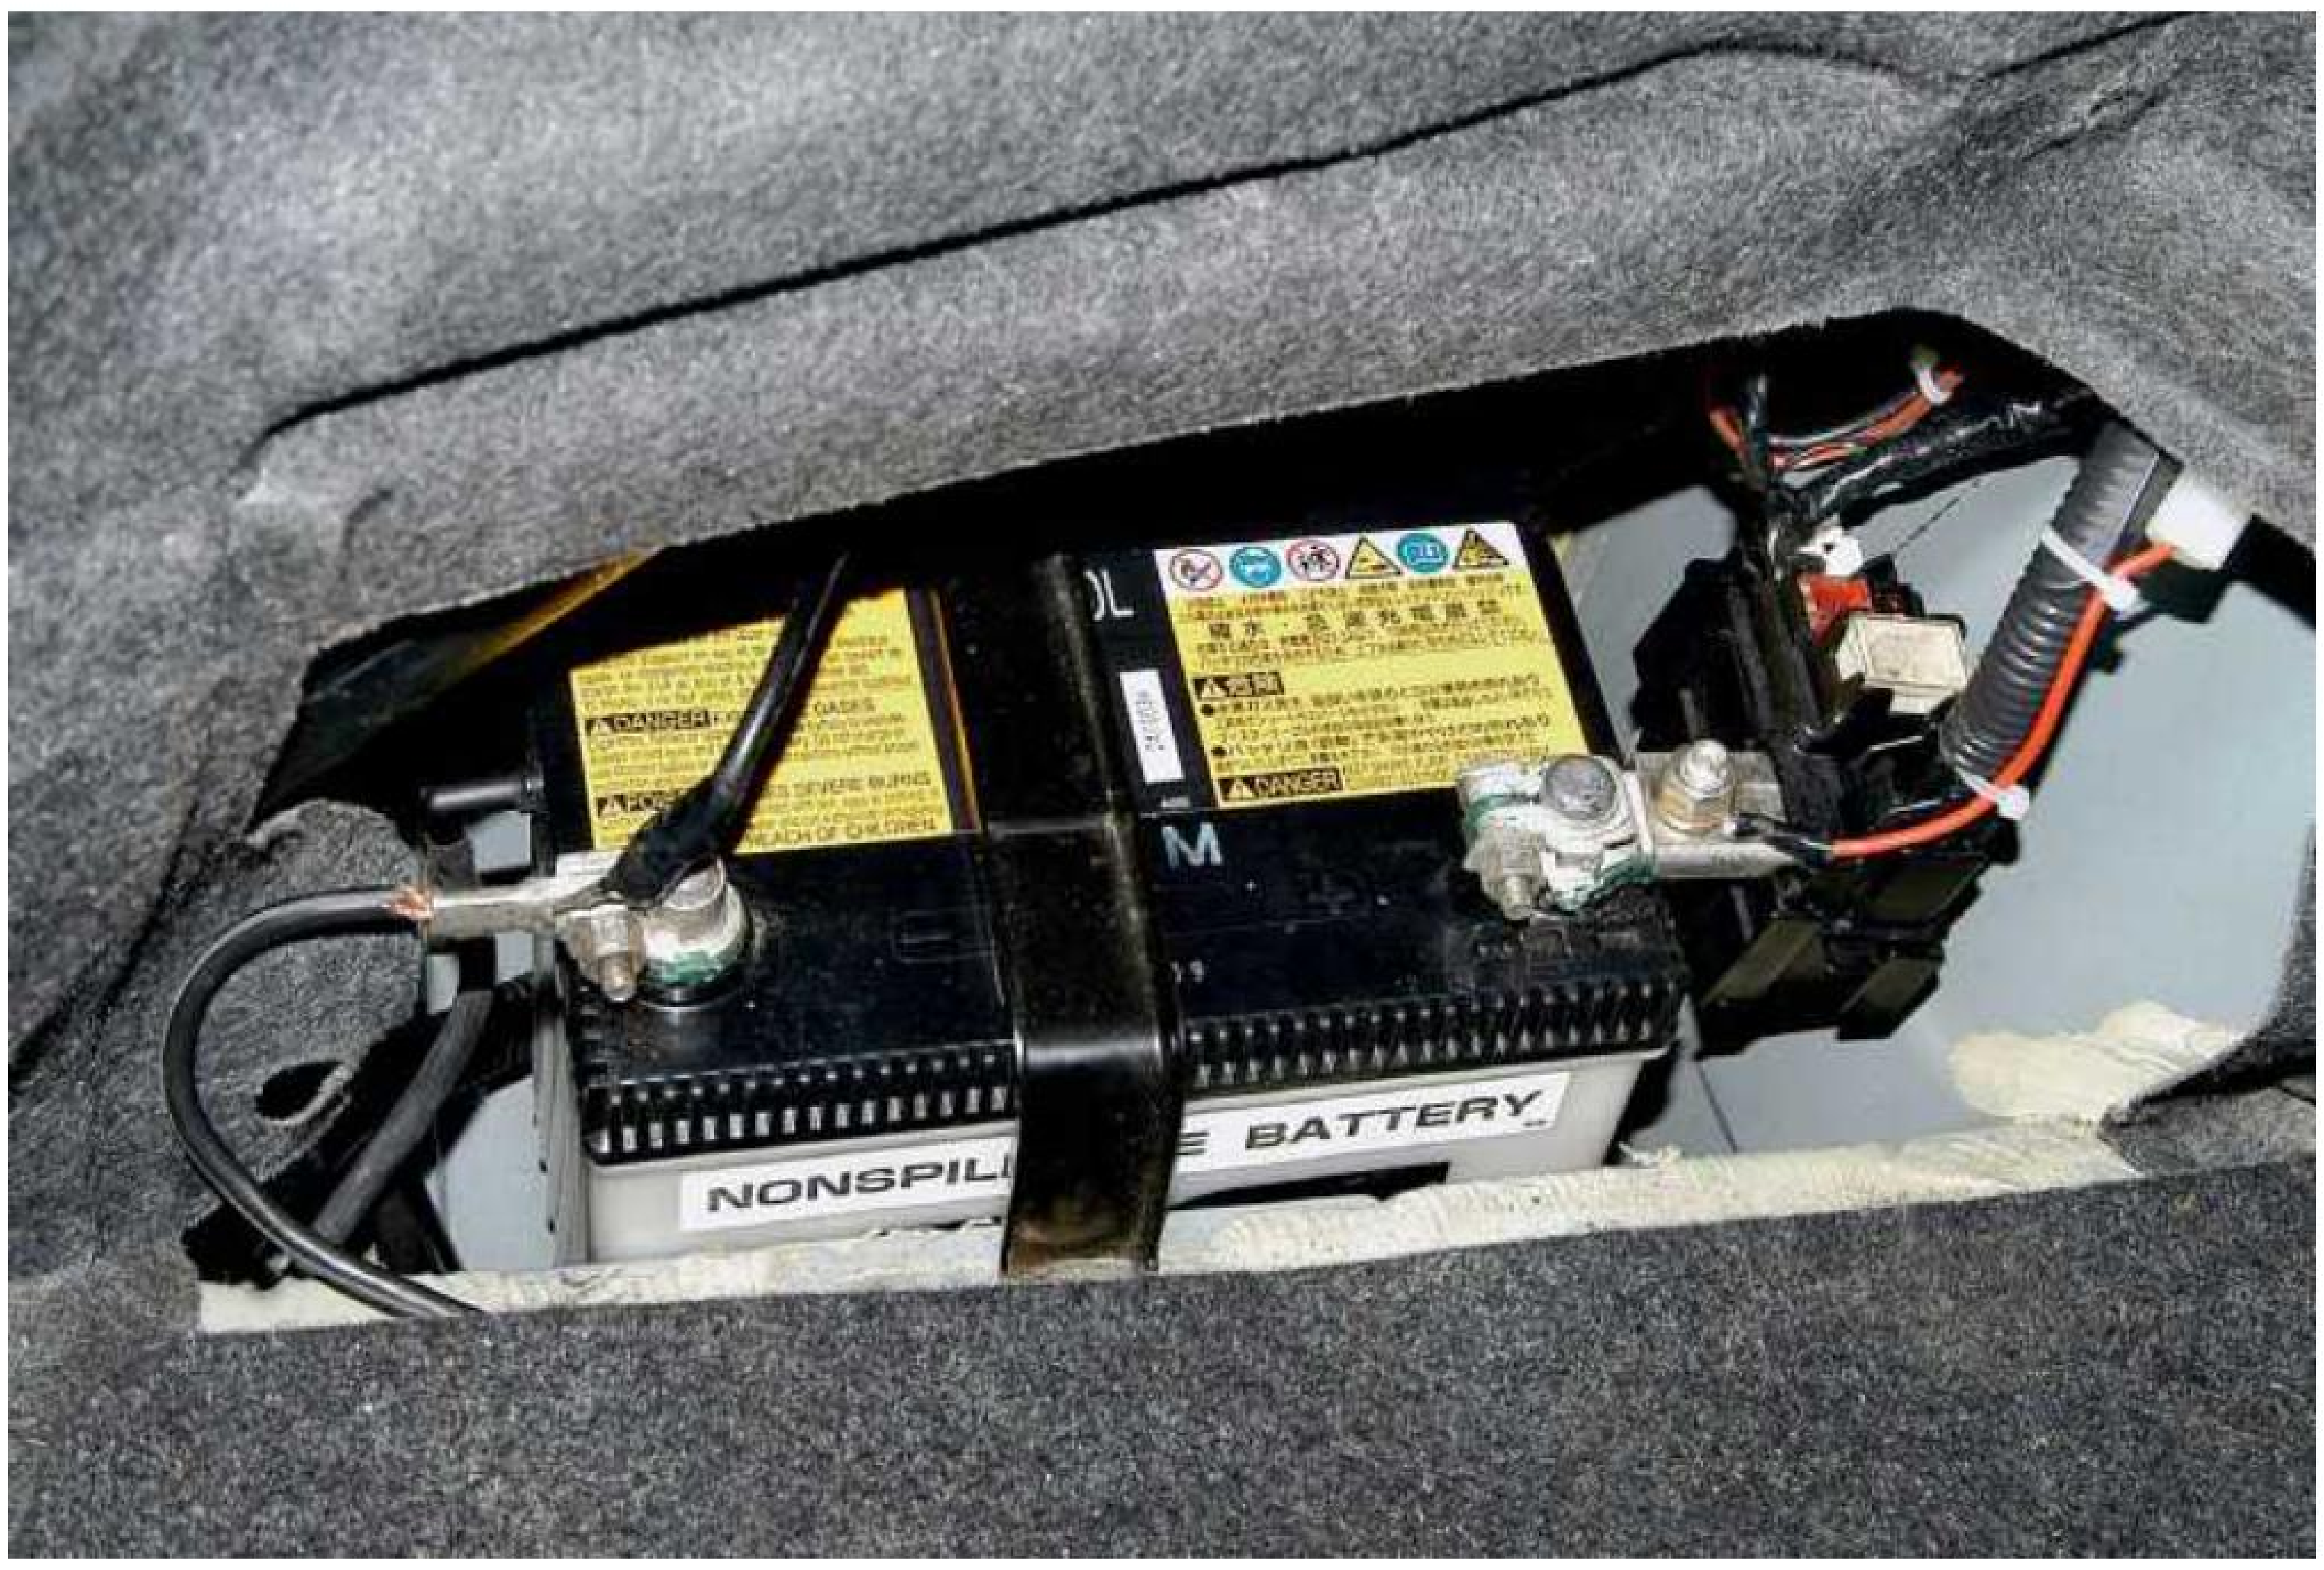

Photo 6: a used car battery that is no longer strong enough to crank a car engine will often be suitable for this application. Such batteries are available free of charge from car workshops and battery replacement shops. However, it's best to get a battery with bolt-on terminals rather than round battery posts like this one. If mounting the battery inside, ensure it is of a sealed design.

First, connect the battery to the solar controller, ensuring the polarity of the connections is correct. Once these connections are made and insulated, insert the inline fuse. The controller should then come alive.

Set the controller to the correct battery type, and if the controller has these facilities, ensure the settings for float charge and auto-disconnect (that occurs if the battery is discharged too far) are correct – many simpler charge controllers won't have these functions.

Next, connect the solar panel to the controller, again ensuring correct polarity of the wiring. If you're unsure of which wire is which (an easy confusion to occur if you've extended the solar panel wiring), use a multimeter to identify the positive and negative wires – the solar panel will need to be exposed to light when you perform this check.

After the solar panel is connected, most controllers will confirm the panel is generating power, either on the LCD screen or by the simple illumination of an onboard LED.

Finally, if you are not using a built-in USB outlet, connect the external USB converter(s). Ensure that no timers are activated on the controller output – you want the output on all the time.

Conclusion

What has surprised me over the four years that I have been running the system is its ease of use and convenience. I just plug in the devices, and they charge – obvious, huh? But they charge purely on solar power, and they charge irrespective of the weather. Also, they charge through blackouts, so there's that safety advantage to consider as well. PE

VOLTAGE
+0.000 V
CR 2032
DURACELL

SmartProbe

Project by Andrew Levido

The SmartProbe is an extremely handy little device for making voltage & continuity measurements. It won't replace your multimeter, but it is designed to be the first piece of test equipment you reach for when debugging or repairing a circuit.

I have made the SmartProbe very small, measuring just 60mm long, 30mm wide and 15mm thick. That's about the size of a box of matches, for anyone old enough to remember one! A probe fine enough for modern surface-mount circuits (or through-hole parts) is fitted to one end of the case, while a short flying ground lead emerges from the other.

It is equipped with a 128 × 64 pixel OLED display and an audio transducer to show voltage measurements and give feedback, respectively.

There are no buttons or switches. The SmartProbe switches itself on when you pick it up, and off again when it senses no movement for a few seconds. You switch between voltage measurement and continuity modes by tapping it with your index finger. The display automatically flips right-way-up however you hold it.

The unit is powered by a single CR2032 coin cell that should last many months with typical use.

Specifications

The SmartProbe operates in either voltage measurement or continuity mode. In the former, it can measure voltages up to ±50V, switching automatically between two ranges. For input voltages below ±6V, it has a resolution of 5mV or better; for higher voltages, the resolution is 50mV.

Its absolute accuracy is within ±0.5%, and the input impedance is around 1MΩ on both ranges.

In the continuity mode, the SmartProbe sources a low current (around 1mA) and displays the voltage drop seen across the probes, just as your multimeter does on the diode test range. The source voltage is 3.3V, enough to forward-bias typical diodes, transistor junctions and most LEDs.

If the resistance between the probes is greater than about 60kΩ, the display shows "OPEN". There is an audio indication of continuity if the measured drop is less than about 1V. The continuity beep responds within a few milliseconds, which is essential for a good user experience.

The SmartProbe is not suitable for use with high-voltage or mains-powered circuits. While it has a degree of input protection, it does not have the insulation or overload ratings of a good multimeter. I considered adding AC voltage or frequency measurement capabilities, but elected to keep it really small & simple.

Features & Specifications

- *Compact size (60 × 30 × 15mm) and lightweight (24g)*
- *Measures voltage or continuity*
- *128×64 pixel OLED screen*
- *Internal buzzer for continuity checking*
- *Measures up to ±50V*
- *Can also test diodes/LEDs and measure forward voltage*
- *Single fine-tipped probe with a ground clip*
- *Powered by an internal CR2032 coin cell*

Ultra-low power design

One of the main aims and challenges of this design was to keep the power consumption when the device is 'off' to a level that would give meaningful battery life, while still being able to sense movement and wake up. A CR2032 cell has a capacity of about 235mAh while discharging from 3V (fully charged) to an end voltage of 2V.

I set myself the goal of aiming for a shelf life of one year (8760 hours), meaning less than 27μA of idle consumption. As we shall soon see, that goal was more than met.

Before diving into the circuit, it is helpful to look at the block diagram of the front end (Fig.1). This shows the device in voltage measurement mode, with just one voltage range for simplicity. We want to convert a bipolar voltage up to ±50V, appearing between the probe and clip, to a unipolar voltage between 0V and 3.3V (V1) suitable for the analog-to-digital converter (ADC).

We do this by fixing the bottom (ground clip) end of the R_A/R_B voltage divider to half the supply rail, ie, around 1.65V. Voltage V1 is therefore

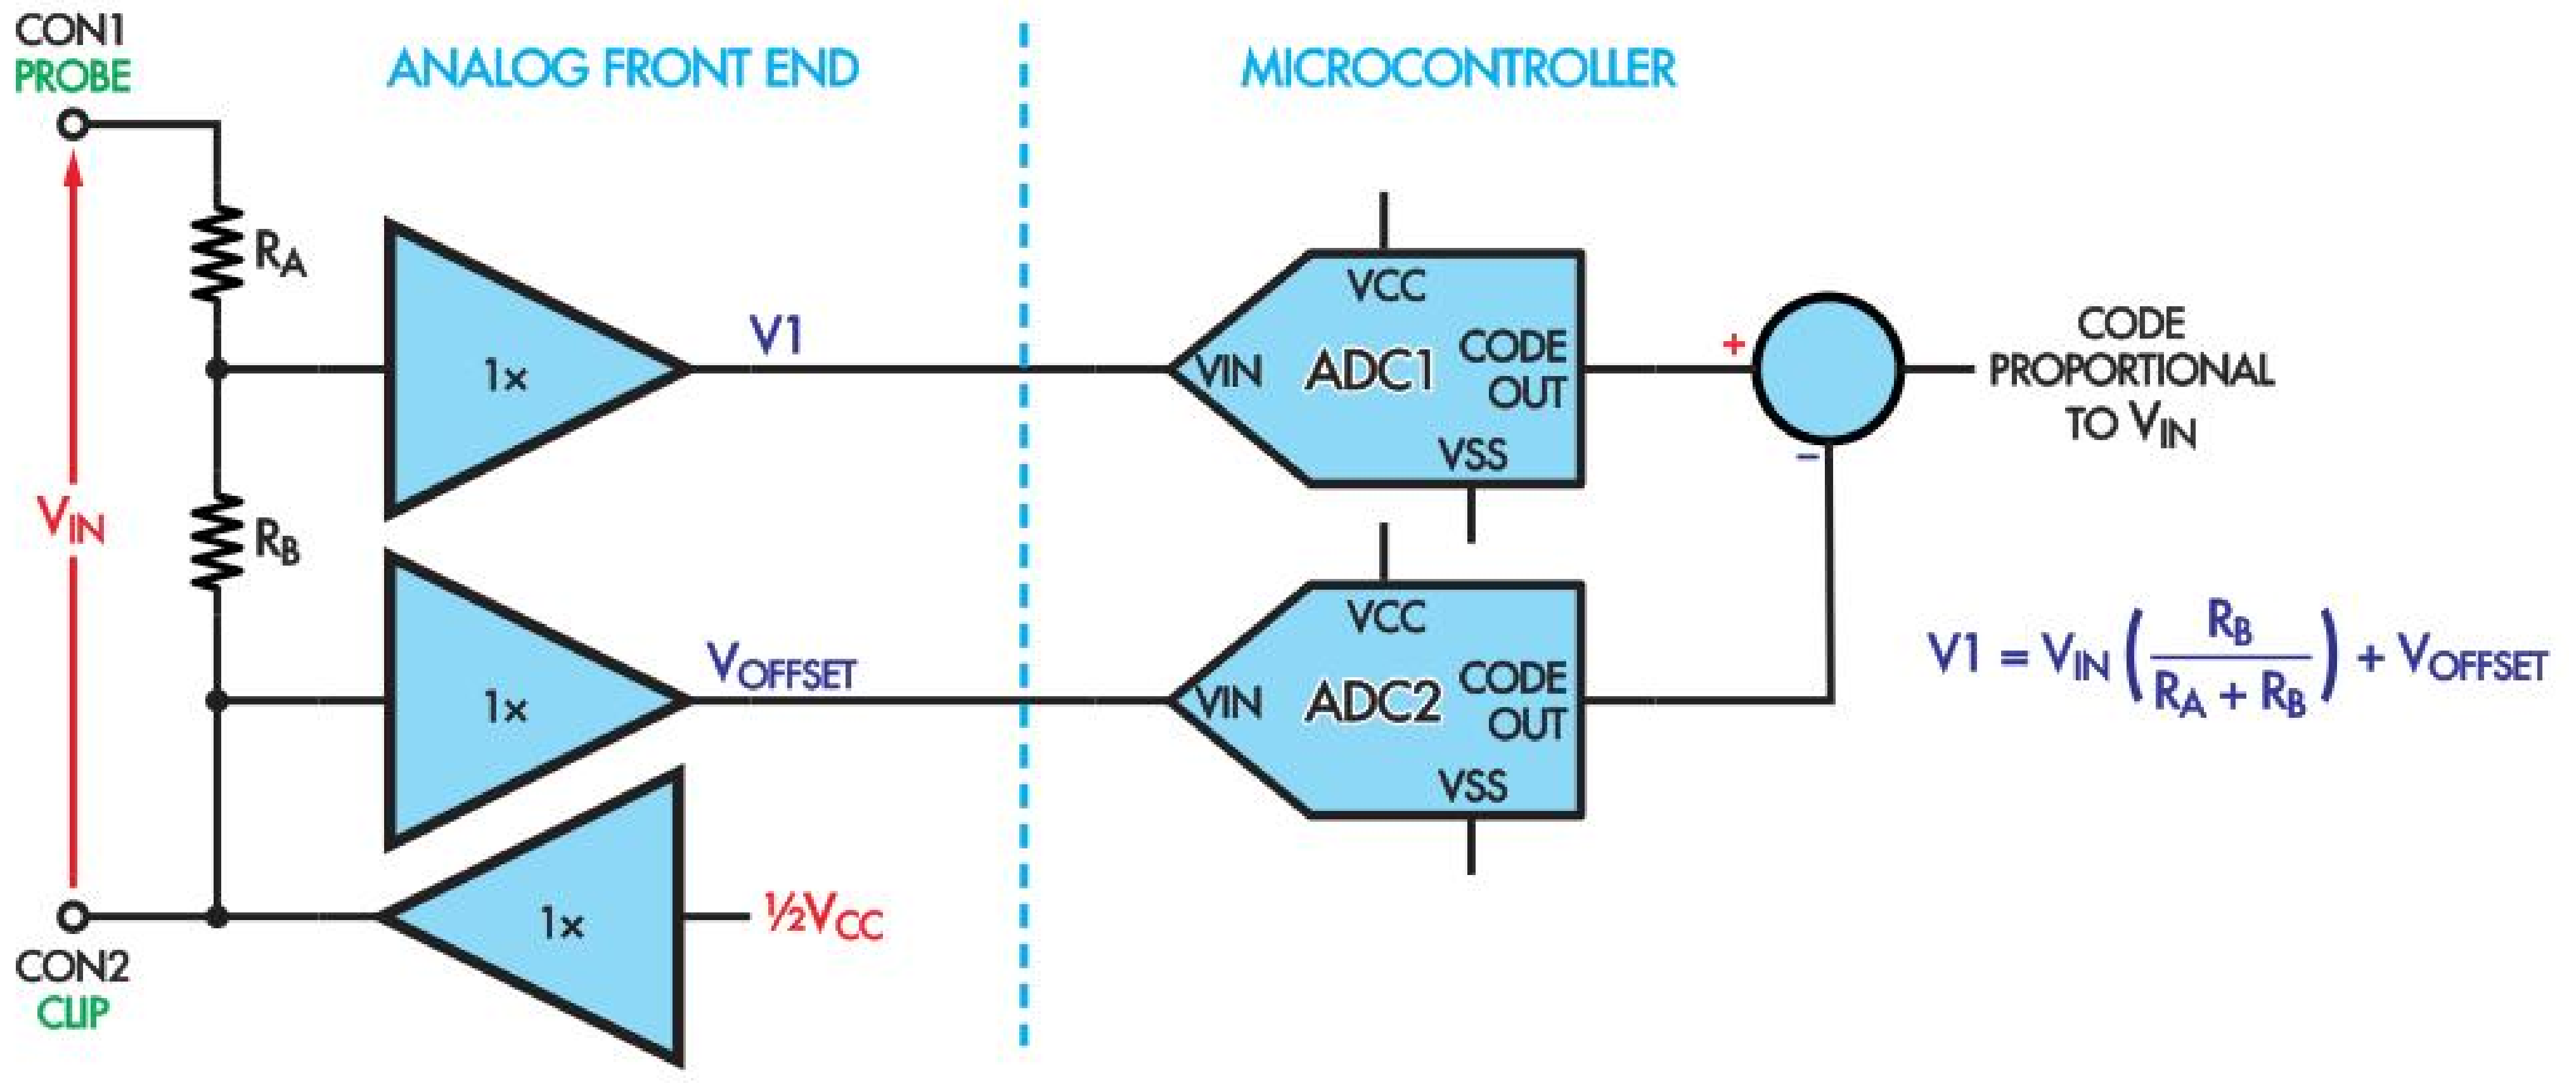

$$V1 = V_{IN}\left(\frac{R_B}{R_A + R_B}\right) + V_{OFFSET}$$

Fig.1: one end of the input voltage divider is fixed to ½ of the supply voltage to provide an offset so that bipolar (±) input voltages can be measured with a unipolar ADC. The offset is later subtracted by the firmware.

an attenuated and buffered version of the input voltage offset by $½V_{CC}$, as described by the equation $V1 = (V_{IN} \times R_B) \div (R_A + R_B) + V_{OFFSET}$.

If we also buffer and convert the offset voltage, V_{OFFSET}, we can subtract it from the converted version of V1 in firmware. The resulting digital code will be a signed value proportional to V_{IN}.

The full circuit (Fig.2) shows that there are actually two dividers and associated buffers in the SmartProbe, one for each input voltage range. For the high-voltage range, it uses a 2MΩ/51kΩ divider, while the 2MΩ/680kΩ divider is for the low-voltage range. The output of each is buffered by IC1d and IC1c, respectively, and fed to its own ADC input channel on microcontroller IC2.

I chose the divider resistor values such that the voltage span seen by the ADC inputs is around 3V centred on around 1.65V (0.15V to 3.15V). This allowed me to stay away from the very ends of the ADC range and avoid a potential source of errors. I am using ±0.1% tolerance resistors here, as these are critical to achieving the required resolution and accuracy.

The op amps are low-cost zero-drift (auto-zero) op amps. These have a worst case offset voltage of ±10µV with just 50nV/°C drift. Since they are connected as unity-gain buffers, there is no appreciable gain error.

The offset voltage is also buffered (by IC1a) and fed to another ADC channel. All three of these input buffers are identical, including protection diode pairs (D1 through D3), and a small amount of low-pass filtering on the inputs (using 1nF/100nF capacitors) and outputs (1.5kΩ/1nF). Most of the filtering of these signals occurs in software, as discussed below.

When the input voltage is outside the ±6V range, the output of the low-voltage sensing circuit (IC1c) saturates, and the digital code associated with this input moves outside the expected range. The firmware automatically switches to using the high-voltage input in this case. If the high-voltage input approaches saturation, a voltage over-range warning message is displayed.

The bottom ends of the dividers are fed by a current-limited buffer, IC1b. The 10Ω resistor provides a bit of protection to the op amp, since its output would otherwise be connected directly to the ground clip and therefore exposed to the outside world. The input of the buffer is connected to a 100kΩ/100kΩ voltage divider fed from one of the microcontroller's GPIO pins (PA06, pin 12).

This pin is configured as a digital output. If it is high, the buffer input is half of the supply voltage, as in Fig.1. If the output is low, the bottom end of the divider is effectively connected to 0V – a state which comes in handy for continuity mode.

Continuity measurement

So far, we have ignored the network consisting of Mosfets Q1-Q3 and the associated passive components. These form an analog switch that is off in voltage mode.

In continuity mode, the offset at the bottom of the voltage dividers is set to zero, as mentioned above, and the analog switch is on. This connects the 3.3V supply to the input probe via the 3kΩ resistor. A voltage of approximately 3.3V is therefore present across the probes when they are open circuit.

This voltage drops as the impedance between the probe falls, ultimately to zero if the probe is shorted to the clip. If a diode junction is connected across the input, with its anode to the probe, the forward drop of the diode will appear across the input.

In continuity mode, the voltage is read by the ADC in the same way as already described, except the offset voltage will be close to zero. The equation shown in Fig.1 will still hold, but we will no longer be able to read negative voltages; that doesn't matter in continuity test mode.

You will notice that the output of buffer IC1c connects directly to the PA03 pin of the microcontroller, as well as to the ADC input via the RC filter. PA03 is internally connected to a fast comparator that drives the continuity beep tone.

The analog switch

The analog switch deserves a closer look. We require a switch with a very high impedance when off, so that no appreciable current flows through the 3kΩ resistor when measuring voltages. The switch must withstand ±50V when open, but have relatively low on-resistance when closed.

I could not find a suitable off-the-shelf analog switch because the voltage requirements are relatively high, so I built my own using two P-channel Mosfets with an N-channel Mosfet to drive them.

The switch is open when Q3 is off, and the gates of Q1 & Q2 are held at their source potential by the 100kΩ resistor. Since the gate-source voltage is zero, both Mosfets will be off. You can see what happens when an external voltage is applied by referring to the left and middle diagrams in Fig.3.

If the drain of Q1 is at +50V, its source will also be at this potential due to the conduction of its body diode. Both Mosfet's gates and sources will therefore be at +50V, so they will remain off. Q2 will therefore block any current flow, since its body diode is reverse-biased.

If the drain of Q1 is at -50V, the body diode of Q2 is forward-biased, leaving the sources and gates of both Mosfets at 3.3V. Q1 now blocks any current flow.

When Q3 is switched on, the gates of both Mosfets gates will be at 0V. The drain of Q2 is fixed at 3.3V, so the body diode initially conducts, causing the sources of both Mosfets to rise almost to this value. The resulting -3.3V gate-source potential is enough to switch both Mosfets on, effectively shorting

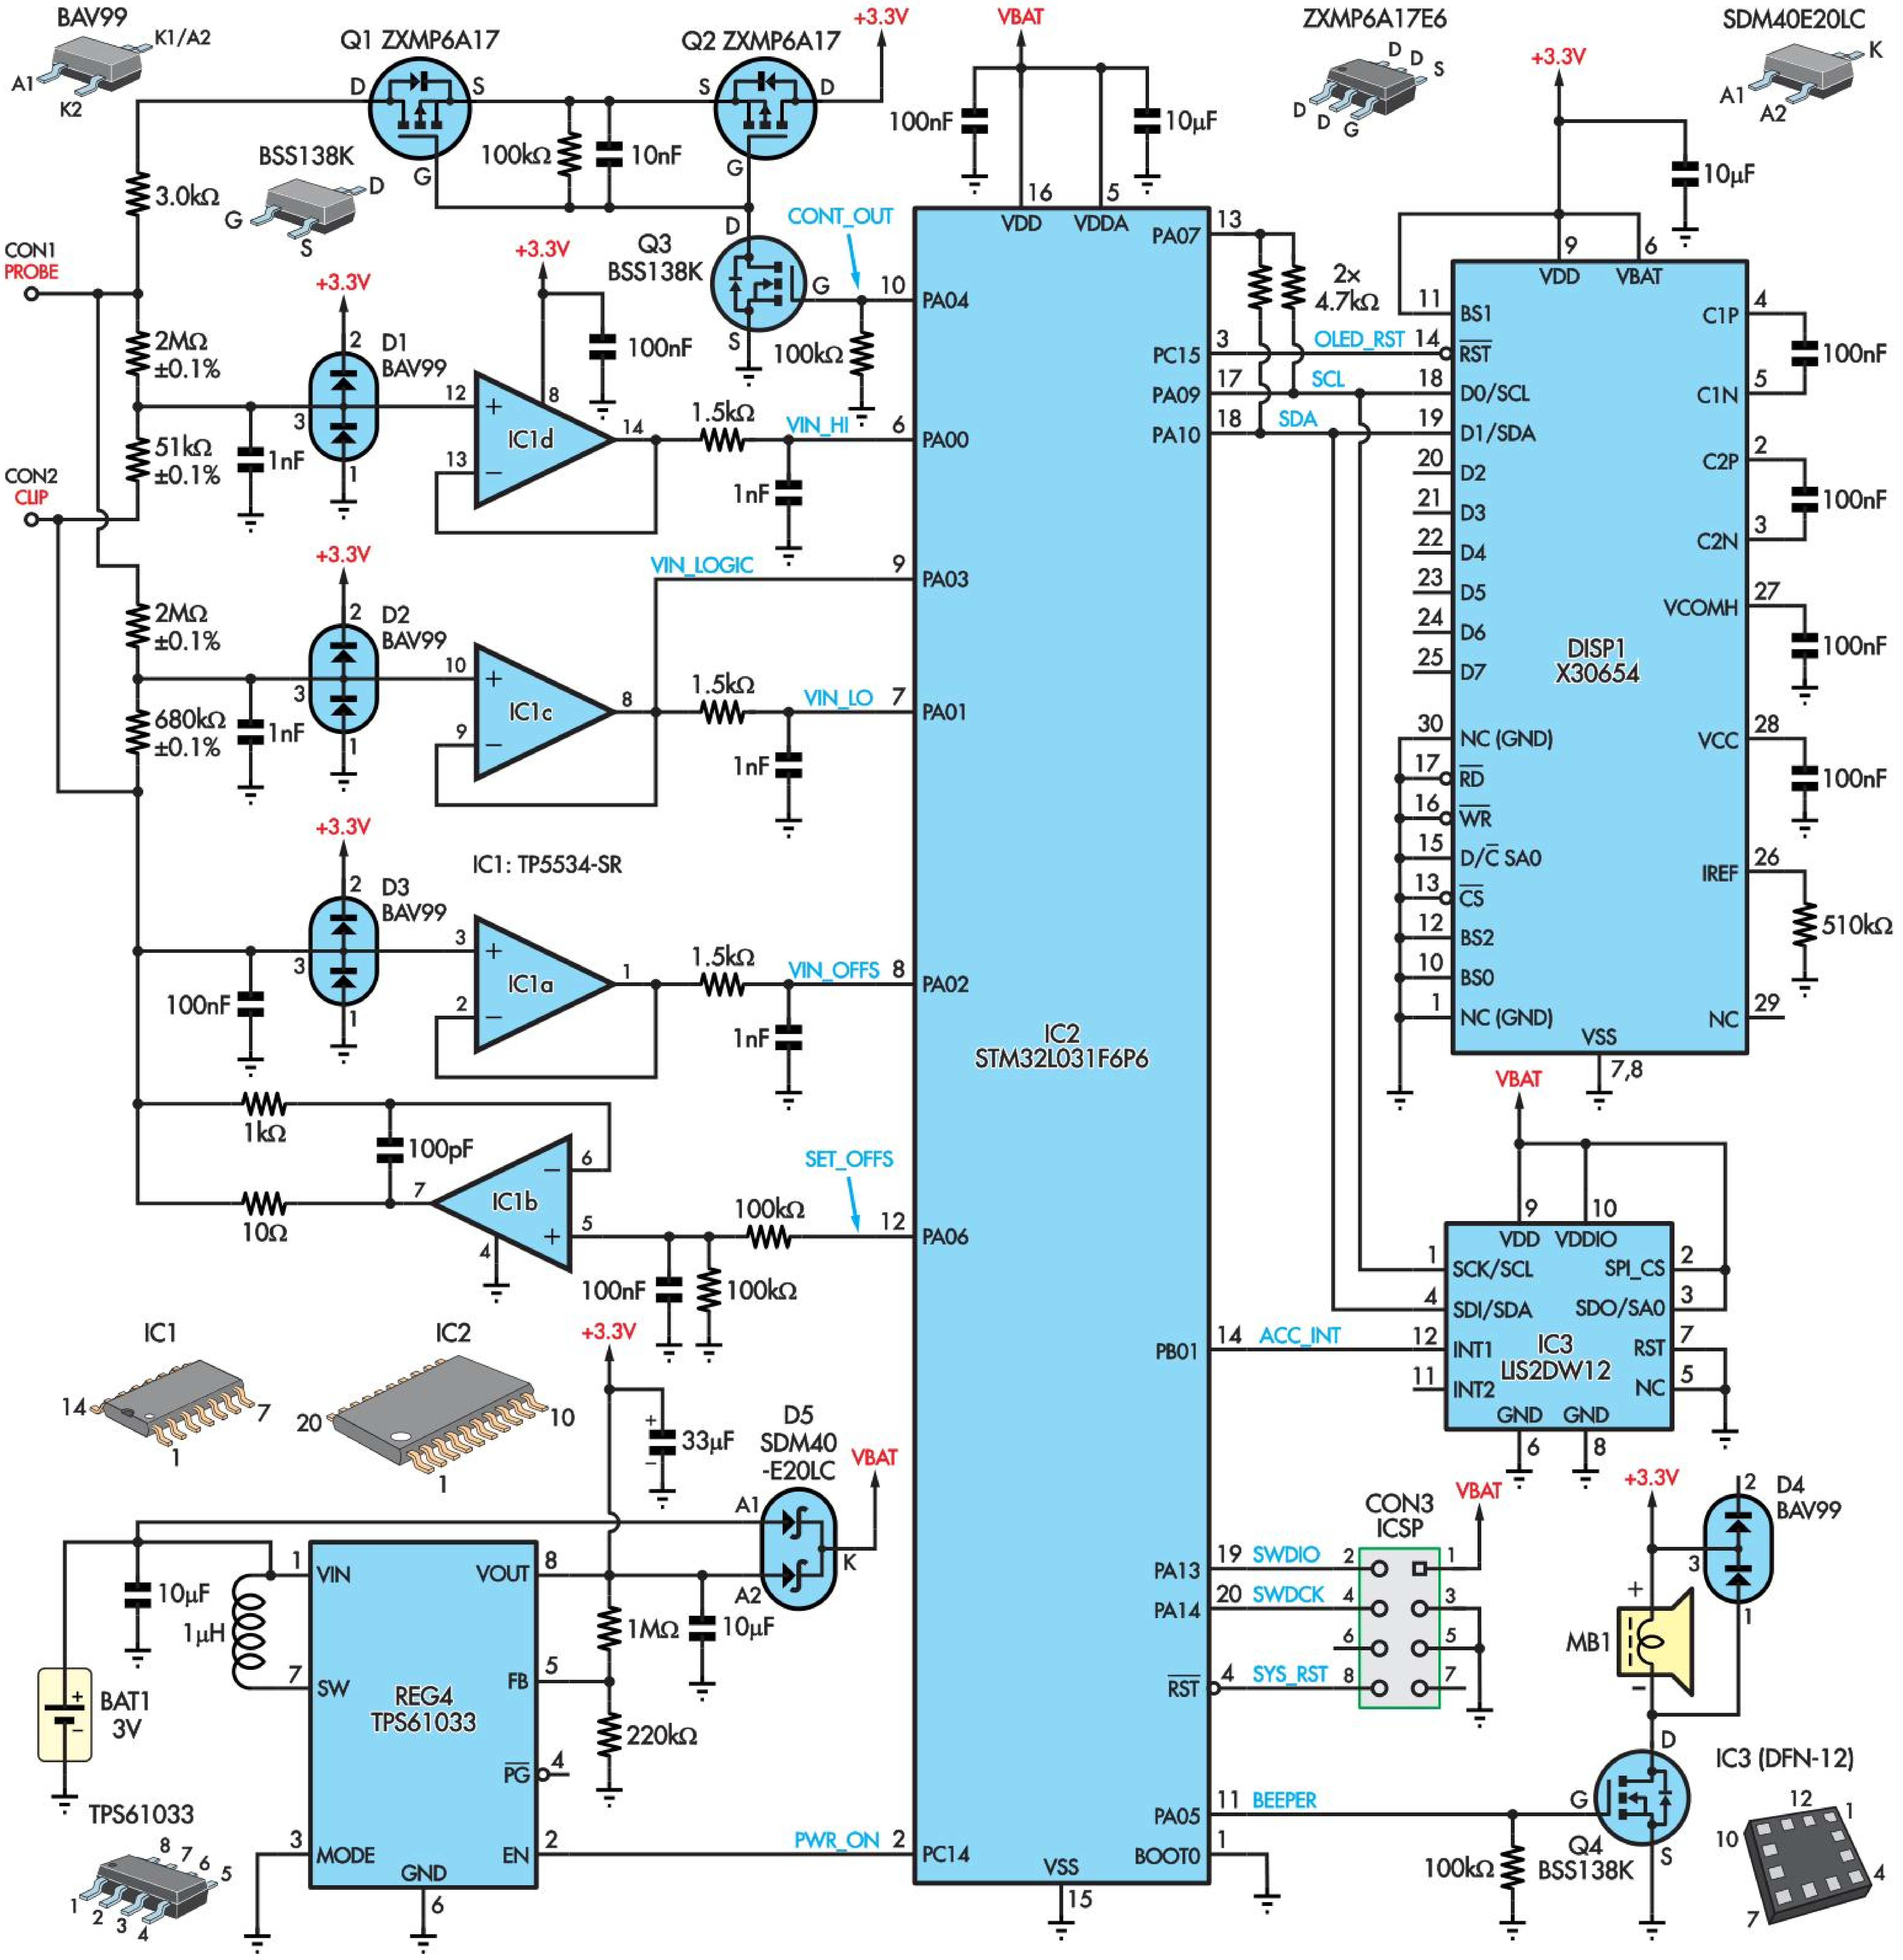
BAV99
K1/A2
A1
K2
Q1 ZXMP6A17
Q2 ZXMP6A17
+3.3V
VBAT
ZXMP6A17E6
SDM40E20LC
BSS138K
3.0kΩ
100kΩ
10nF
100nF
10μF
CON1 PROBE
CON2 CLIP
2MΩ ±0.1%
51kΩ ±0.1%
680kΩ ±0.1%
1nF
D1 BAV99
D2 BAV99
D3 BAV99
IC1d
IC1c
IC1a
IC1b
IC1: TP5534-SR
Q3 BSS138K
CONT_OUT
1.5kΩ
VIN_HI
VIN_LOGIC
VIN_LO
VIN_OFFS
SET_OFFS
1kΩ
10Ω
100pF
IC2 STM32L031F6P6
VDD
VDDA
PA07
PA04
PC15
PA09
PA10
PA00
PA03
PA01
PA02
PA06
PB01
PA13
PA14
RST
PA05
PC14
VSS
BOOT0
2× 4.7kΩ
OLED_RST
SCL
SDA
DISP1 X30654
BS1
D0/SCL
D1/SDA
D2
D3
D4
D5
D6
D7
NC (GND)
RD
WR
D/C SA0
CS
BS2
BS0
VCOMH
VCC
IREF
C1P
C1N
C2P
C2N
NC
510kΩ
IC3 LIS2DW12
SCK/SCL
SDI/SDA
INT1
INT2
SPI_CS
SDO/SA0
VDDIO
GND
ACC_INT
CON3 ICSP
SWDIO
SWDCK
SYS_RST
BEEPER
PWR_ON
IC1
IC2
33μF
D5 SDM40 -E20LC
1MΩ
220kΩ
1μH
BAT1 3V
REG4 TPS61033
VIN
VOUT
SW
FB
PG
MODE
EN
TPS61033
D4 BAV99
MB1
Q4 BSS138K
IC3 (DFN-12)
SC ©2025 SMARTPROBE

Fig.2: the SmartProbe circuit is fairly straightforward, except for a few tricks related to achieving ultra-low power consumption that are detailed in the text.

out their respective body diodes and switching the analog switch fully on.

The ZXMP6A17E6 Mosfets I chose have a maximum V_{DS} of -60V and an $R_{DS(ON)}$ of less than 0.5Ω with a V_{GS} of 2.5V (interestingly, they have six pins, but the three additional ones are just extra drain connections). The total on-resistance of the analog switch will therefore be around 1Ω. The P-channel Mosfets see a worst-case voltage of -53.3V.

The BSS138K (Q3), with a maximum V_{DS} of 50V, places the upper limit on switch voltage. This is a bit of a soft limit, since the 100kΩ resistor limits the avalanche current if Q3 were to break down. Nevertheless, this switch is what limits the nominal maximum voltage for the SmartProbe.

Digital circuity

The microcontroller is a 32-bit STM32L031F6 low-power Arm Cortex M0+ from ST Microelectronics. This has 32kiB of flash, 8kiB of RAM and comes in a 20-pin TSSOP package. Importantly, it operates at any voltage between 1.8V and 3.6V, and can be put into various low-power modes where its current draw reduces to single-digit microamp levels while still retaining RAM contents.

The display connects to the microcontroller via an I^2C interface. This is a 128 × 64 pixel white OLED screen that

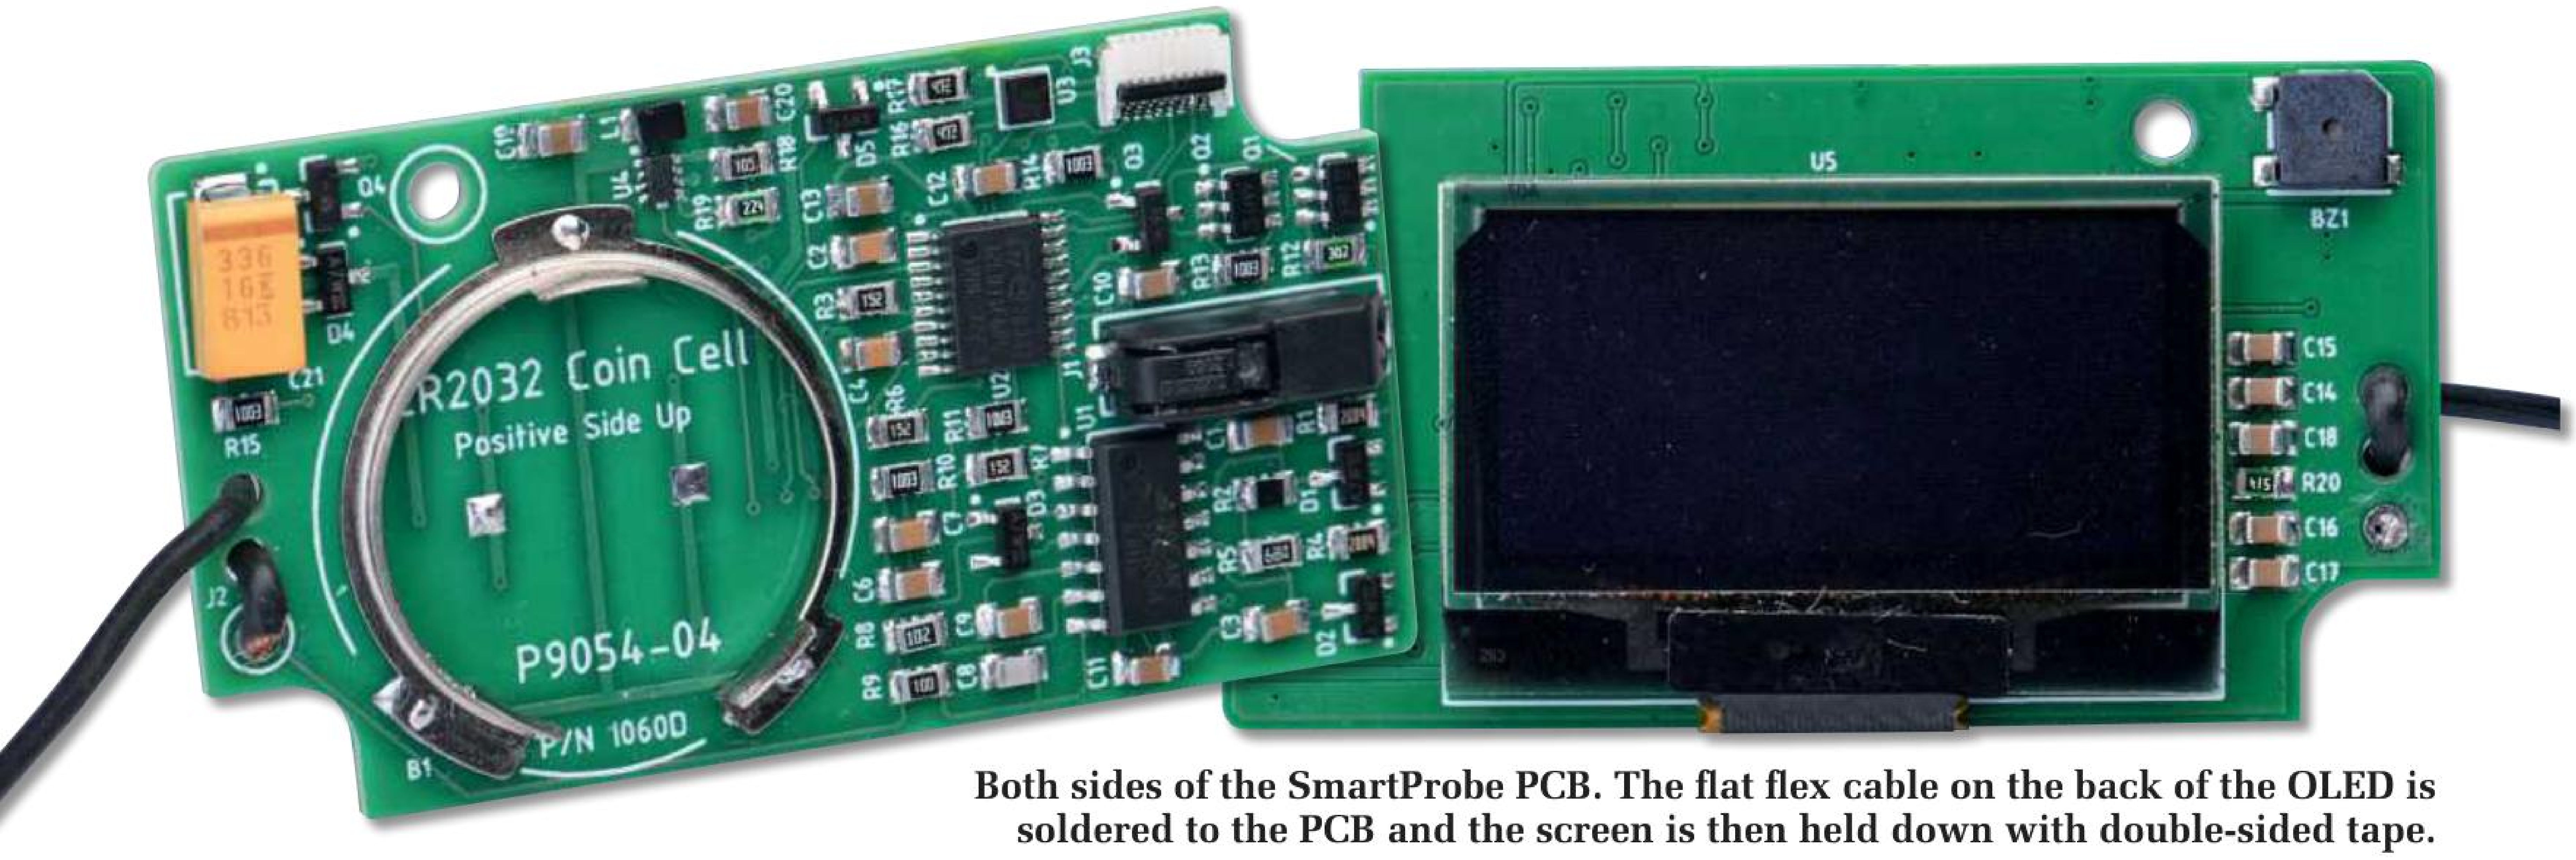

Both sides of the SmartProbe PCB. The flat flex cable on the back of the OLED is soldered to the PCB and the screen is then held down with double-sided tape.

measures just 34 × 22 × 1.5mm. These displays are readily available for just a few dollars each from AliExpress. They have a SH1106 control chip onboard, and can be controlled via an 8-bit parallel or SPI/I²C serial bus.

The display only needs four 100nF capacitors and one resistor added to create the necessary internal voltages (all shown to the right of DISP1).

The OLED current is set by the resistor. I have used 510kΩ, which gives a current of about 10μA per pixel. This is nice and bright, but means that the display could draw as much as 82mA (128 × 64 × 10μA) if all pixels were on. In reality, the measured current stays under 20mA or so, since we never light more than about 25% of the pixels simultaneously.

Nevertheless, the display is the main current consumer in the circuit and has to be completely shut down when the SmartProbe is inactive. To help ensure it comes up reliably when awakened, I have wired its hardware reset pin to a microcontroller GPIO pin (PC15).

The audio transducer, MB1, is a small magnetic beeper that requires an AC signal to operate. This means it can deliver a variety of tones if driven at different frequencies. I have used one of the micro's PWM outputs (routed to the PA05 pin) to drive it via Mosfet Q4.

Using a PWM output allows me to set the tone by varying the PWM carrier frequency and manage the volume, and more critically, the current consumption, by limiting the duty cycle. Even so, it presents a significant (relatively speaking) load on the supply. For this reason, I used it only when I think it is really necessary.

Accelerometer

Because the accelerometer is the only component in this circuit that is always operational, I needed to choose it carefully. I selected the LIS2DW12 three-axis "digital motion sensor" for this application because it is specifically geared toward ultra-low-power applications.

This chip (IC3) contains a three-axis MEMS accelerometer and a heap of signal processing hardware that can be configured to detect device orientation, free-fall events and tap or double-tap events on any axis. It can detect activity and put itself into a low-power state when it senses inactivity, waking itself up again autonomously.

I took advantage of the orientation function to flip the display the right way up, the single tap function to change modes and the activity/inactivity function to turn the SmartProbe on and off.

The device supports a range of sampling rates and draws anywhere between 500nA and 90μA when operating. Lower data rates result in lower operating currents, but some features we need won't work at the very lowest data rates. For this reason, we run the accelerometer at 400 samples per second when active, dropping to 200 samples per second when 'off'.

This gives us an 'off' power consumption of somewhere around 12-20μA. The data sheet says the consumption will be 12μA in the mode and data rate we use, but that is specified at 1.8V and 25°C. The data provides no help in understanding what the consumption will be at higher temperatures or with voltages up to 3.0V, as it will experience in our circuit.

No data usually means you can safely assume it will be worse. My measurements show consumption closer to 16-20μA with our battery voltage range and my (unairconditioned) room temperature.

While that is higher than the published figures, I think it is still pretty amazing performance considering the chip is taking three 14-bit accelerometer samples every 5ms and pushing them through a fairly complex digital signal processing chain.

Hardware-wise, the LIS2DW12 is very nice; it has just 12 pins, requires no external components and costs just $2.50 in single quantities. The downside is that it is only available in a 2 × 2mm leadless package. Fortunately, it does not have a thermal pad, so it is easier to hand-solder than some chips I have come across.

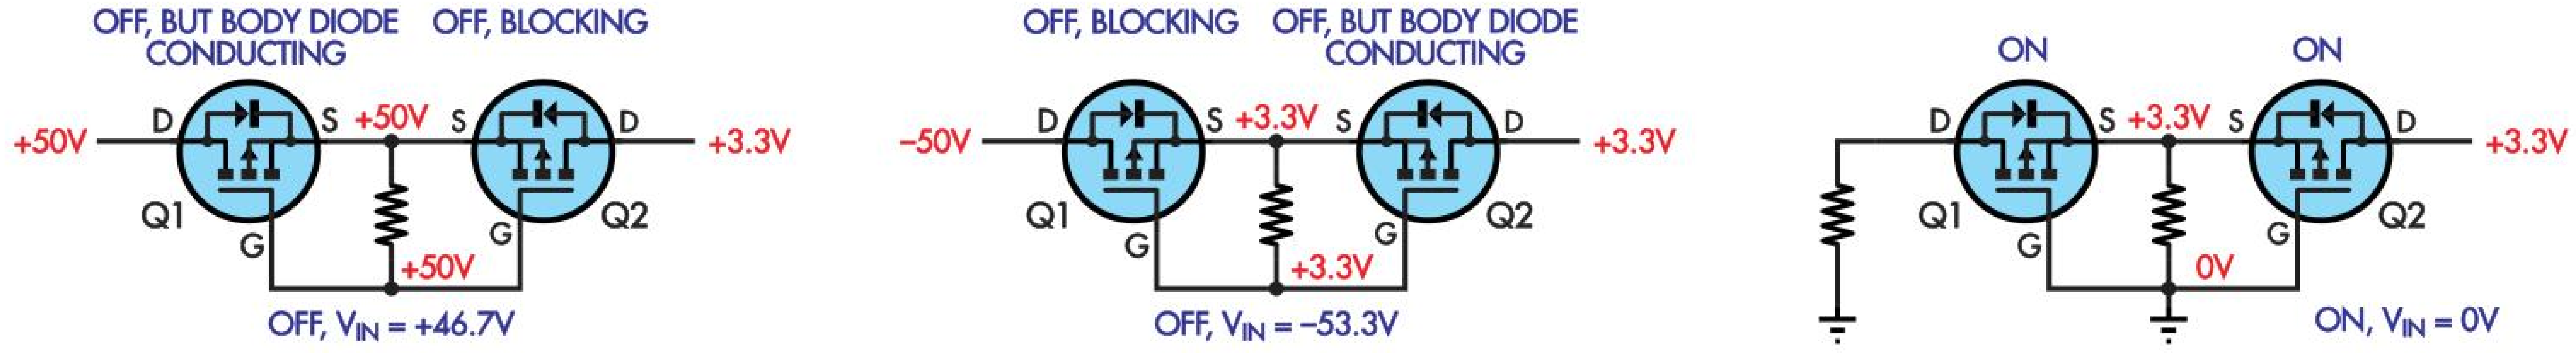

Fig.3: the left and middle diagrams show the voltages on the analog switch in the off state with +50V and -50V applied to the input, respectively. The rightmost diagram shows them when the switch is on.

It is a complex device from a firmware perspective, but I took the time to write a (reasonably) comprehensive driver, since I intend to use this chip again.

Power supply

The power supply scheme is straightforward, as shown in Fig.4. There are two power rails: V_{BAT}, which is derived from a lithium coin cell and is always available, plus a 3.3V rail that is only available when the SmartProbe is on. The 3.3V rail is derived from the coin cell via a boost converter based around IC4, a TPS61033.

The boost converter is enabled by the PWR_ON signal from the microcontroller, so the display, analog front end and beeper are only powered when the SmartProbe is on. The shutdown leakage current of the TPS61033 is specified at 0.1µA, so well within our meagre power budget.

The V_{BAT} supply is the diode-OR combination of the coin cell voltage and the 3.3V supply. The upshot of this is that V_{BAT} will be approximately 3.3V while the device is on, but will fall to very near the battery voltage when asleep. Both the accelerometer and microcontroller will happily operate at any voltage between 3.6V and 1.8V, so this is not a problem for them.

I used a pair of schottky diodes to combine the supplies to minimise the forward drop, keeping it to only about 0.2V at the low currents drawn in standby mode.

Firmware

The firmware architecture is shown in Fig.5. The software consists of a main loop, shown at the top, and four asynchronous tasks triggered by interrupts, shown at the bottom. Some of these asynchronous tasks communicate with the main loop via a few shared data registers and flags.

When an interrupt occurs, the processor stops what it is doing and starts running code from the appropriate interrupt service routine (ISR).

When the code starts from reset, the microcontroller core and onboard peripherals are initialised. This includes such things as setting up the microcontroller clocks, the I^2C peripheral, PWM, timers and the like. This only needs to be done once, because when the microcontroller is stopped, the RAM contents and register data are retained.

After this, the external peripherals are initialised via their drivers. This involves enabling the boost converter, initialising the accelerometer and the OLED display. This step is repeated each time the SmartProbe awakens because the display driver's configuration registers are lost when the 3.3V rail is disabled.

We also reconfigure the accelerometer, even though it is never shut down completely. We do, however, disable some of its functions before putting the microcontroller into sleep mode, so it is easier just to reprogram it completely when it wakes up.

Once everything is initialised, we enter the main loop proper. Here, we check if the accelerometer has detected a period of inactivity and put itself into low-power mode. If it has, we proceed to put the SmartProbe into 'off' mode, as described below.

If it is still active, we check if fresh data is available from the ADC sampling task. If not, we loop back and repeat the cycle, checking continuously for inactivity and new data.

When fresh ADC data is available (approximately twice per second), the display is updated according to the operating mode and taking into account the orientation of the SmartProbe. This routine also takes care of the auto-ranging and over-voltage detection.

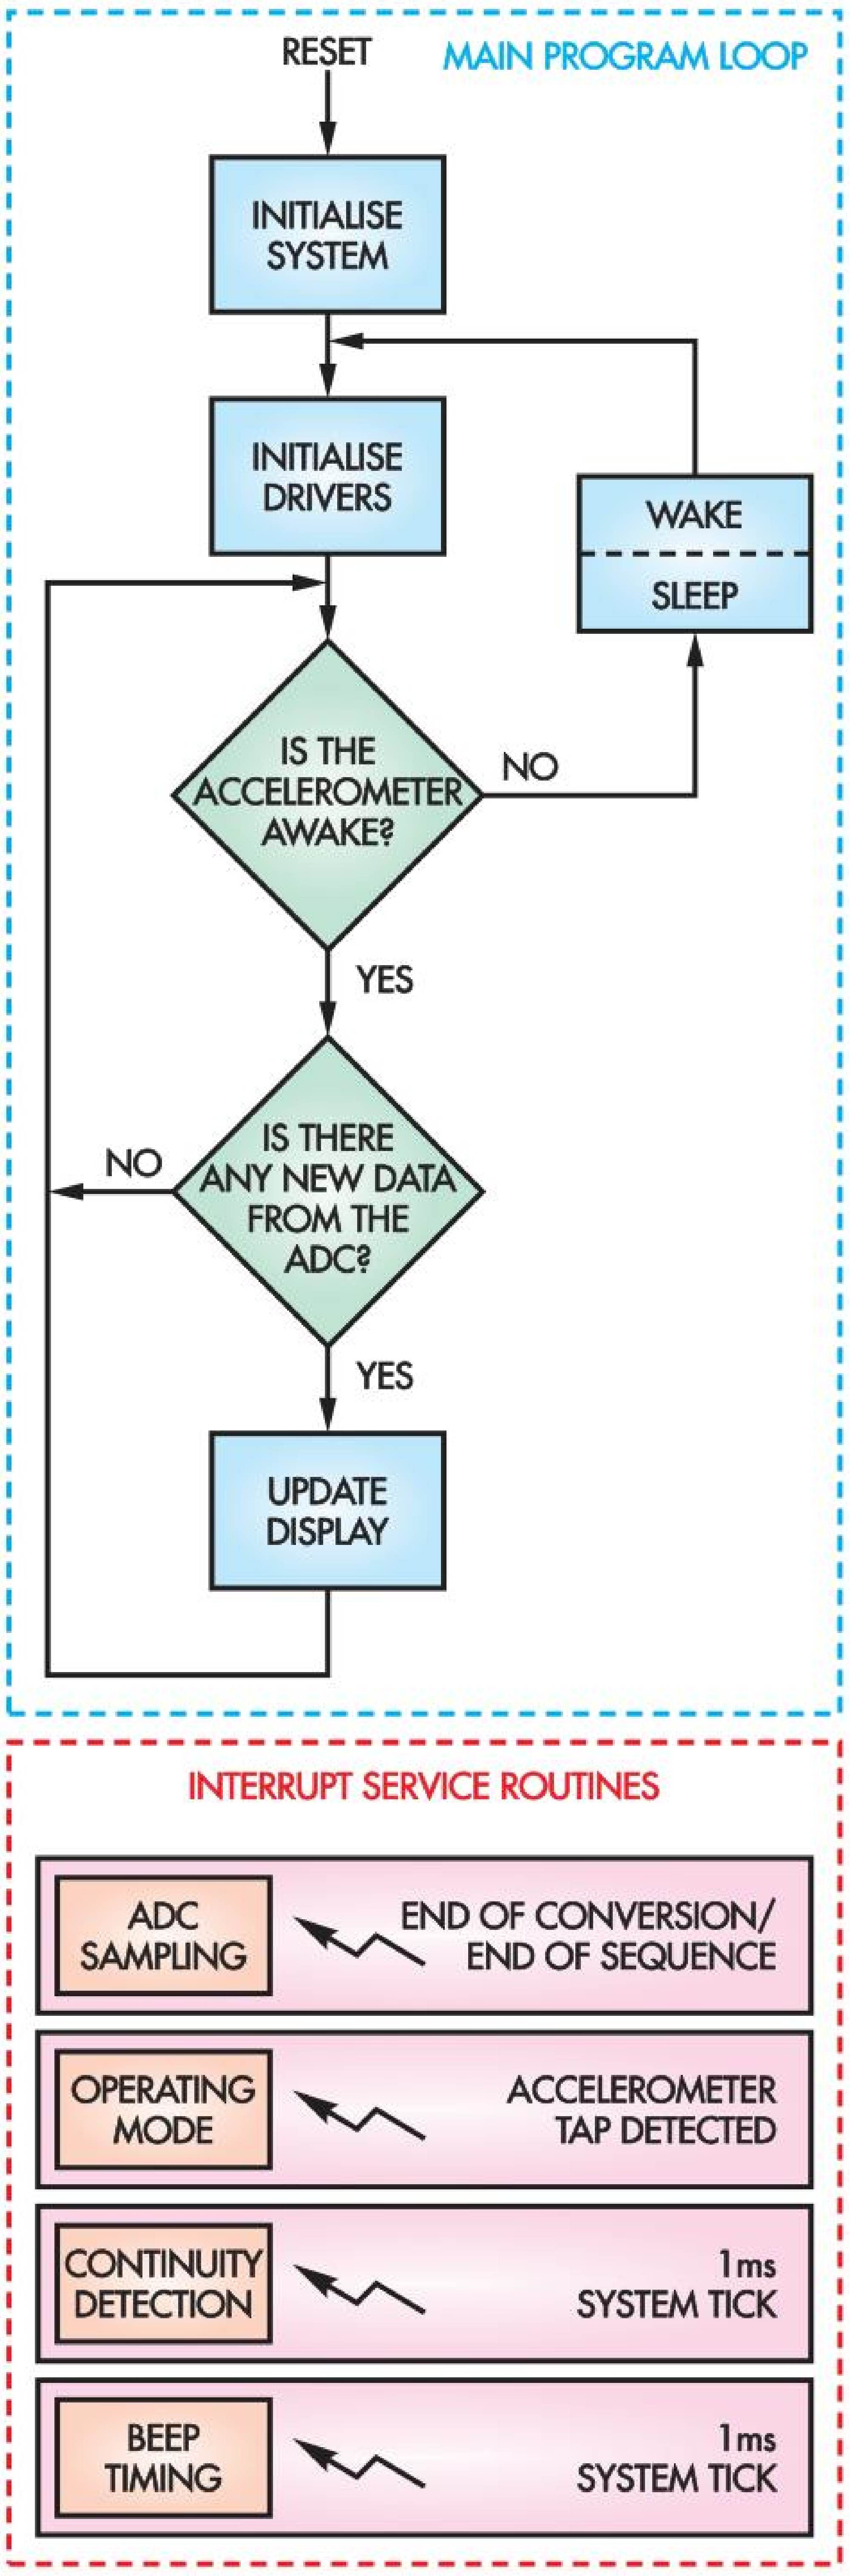

Fig.5: the firmware consists of a main loop and four asynchronous tasks, as described in the text. They communicate with each other via a few shared data and status registers (not shown).

ADC sampling

The ADC sampling routine operates independently of the main loop, triggered every 500ms by a repeating timer. The analog-to-digital converter (ADC) peripheral within this microcontroller is extremely flexible. It is a 12-bit successive approximation converter with an input multiplexer that can handle up to 18 input channels.

Sixteen of the channels can be connected to input pins (not all of which are available on the 20-pin version of the chip), while two can be connected to internal sources. One of these is a 1.2V bandgap voltage reference. The ADC also includes a zero-calibration feature, which we use each time the SmartProbe becomes active.

The ADC can be configured to

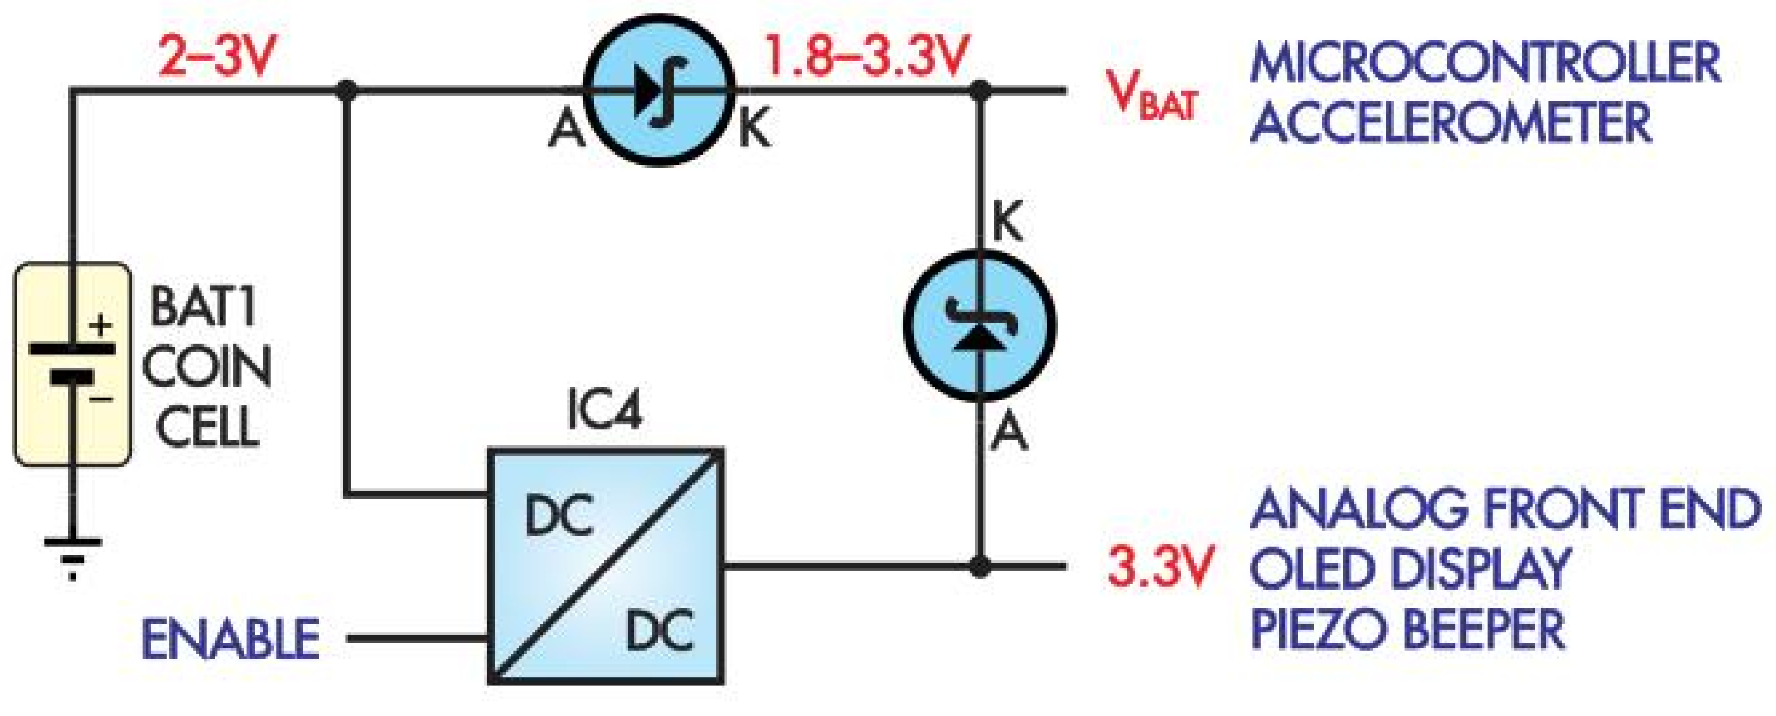

Fig.4: the 3.3V power rail is derived from the coin cell voltage via a boost converter. This is disabled when the SmartProbe is 'off' but the microcontroller and accelerometer remain powered by the V_{BAT} rail.

Table 1 – STM32L031F6P6 power saving modes

Power Mode	Core	Peripherals	RAM/Registers	Wake-up
Run	On	On	Retained	Not Applicable
Sleep	Off	On	Retained	Any Peripheral
Stop	Off	Off	Retained	External interrupts, low-power peripherals, RTC, watchdog
Standby	Off	Off	Lost	RTC, watchdog, wake-up pin

automatically scan and convert channels in a sequence. It also has a hardware oversampling capability. If this is enabled, the ADC will take a series of samples and average the results for each channel in the sequence.

I used these features to create a sampling regime that eliminates a lot of the mains-frequency interference that would otherwise make achieving stable readings difficult.

If we take many samples of a signal over an integral number of mains cycles and average them, any mains frequency component will average to zero. This is because the average of any sinusoidal signal over a full cycle is zero. So, if we make the ADC sample and average each input over one or more 20ms intervals, any 50Hz component will be eliminated.

In the SmartProbe, the clock division options available to us mean that we can't quite do this perfectly. The best we can manage is to take the 256 samples of each input over a period of 59.05ms – very close to three mains cycles. This means the mains cancellation will not be perfect, but we should still reduce it by 34dB (50 times) or thereabouts, which helps a lot.

This oversampling and averaging also serves as a simple low-pass filter, helping to smooth out any small perturbations in the voltage being measured.

The ADC is therefore set up to convert four inputs in sequence: the high and low range voltage measurement inputs, the offset voltage and the internal reference. Each is sampled 256 times, and the results averaged twice per second. Once configured, all this happens more-or-less automatically.

An interrupt is triggered at the end of each conversion sequence, at which point we need to translate the averaged ADC readings into absolute voltages that we can display. The ADC output is ratiometric with the V_{BAT} power rail – that means its output code is a measure of the input voltage as a fraction of V_{BAT}.

V_{BAT} is nominally 3.3V when the SmartProbe is on, but this voltage is not regulated to the extent that would allow conversion to absolute voltages at the level of accuracy we want.

Fortunately, there is a way around this, using the internal bandgap reference. This has a nominal output of 1.2V and pretty good stability. When the chip is manufactured, the value of this internal reference is measured by the ADC while the supply voltage is fixed at a fairly precise 3.00V, and the resulting code burned into non-volatile memory.

The firmware uses this code, together with the real-time measurement of the internal reference, to calculate the V_{BAT} voltage on each measurement cycle.

Knowing the supply voltage allows us to determine the absolute value of the input voltages. It is then only a matter of subtracting the offset voltage from each of the input voltages to determine the voltage across the lower resistor in each divider as a signed integer in units of millivolts. These are later adjusted for the divider attenuation in the display update routine.

When new values are calculated, they are stored in shared memory, and the flag set to let the main loop know that new data is available.

Continuity mode & beeper

The accelerometer is set up to detect single-tap events in the 'vertical' axis of the smart probe (ie, the top or bottom surface when looking at the display) and assert the interrupt line. The associated interrupt service routine responds by switching modes: opening or closing the analog switch and setting the offset voltage appropriately.

A short tone sounds when changing modes – a higher frequency when switching to voltage mode, and a lower frequency when switching to continuity mode.

The beeper driver makes use of the M0+ core's dedicated tick timer, which is usually set to provide a system tick interrupt every millisecond. A tone is initiated by calling a driver function specifying the desired frequency and duration. The function starts the tone playing at the appropriate frequency, then returns. The tone is automatically terminated after the requisite number of 1ms ticks elapse.

It is important (to me anyway) that the continuity beeper has a very fast response. The ADC samples are only updated every 500ms, which is fine for the display, but way too slow for the beep. Therefore, I used a comparator, as mentioned earlier.

One of the two onboard comparators is configured to compare the low-range input voltage with a fixed internal voltage set to ¼ of the internal 1.2V reference. If the voltage at the input pin falls below 0.6V, a flag is set to indicate continuity. If it is above this level, the flag is cleared. The comparator output flag is sampled every system tick and the continuity beep is sounded if it is set.

Low power operation

We mentioned above that the accelerometer is configured to detect a period of inactivity and autonomously put itself into a low-power mode. When the microcontroller detects this has occurred, the rest of the circuit must be shut down until the accelerometer indicates activity has resumed.

We have seen that the accelerometer consumes up to 20μA in its low-power mode, and we have a total design target consumption of 27μA or less. This leaves us with just a few microamps for everything else, including the microcontroller.

The Cortex M0+ architecture supports a variety of low-power modes with differing levels of power consumption. The trade-off for lower power is longer wake-up times and more limited wake-up functionality.

Table 1 shows a (very) simplified chart of the available modes and their key differences. Every one of the modes shown in the table has several variations, and it is entirely possible for a lower power mode to consume more than a higher power mode depending on the exact configuration.

"Run" is the normal operating mode of the microcontroller and has the maximum power consumption; the core and all peripherals are operating. The easiest way to reduce power in this mode (or any mode where the clocks are operating) is to reduce the frequency of the system clock below its maximum of 32MHz. I use a 3MHz system clock in the SmartProbe for this reason.

In "Sleep" mode, the core clock is disabled, but the peripherals remain fully operational. RAM and register contents are preserved. This allows for a very fast wake-up (in the order of 0.35µs) but comes at the cost of around 1mA current consumption at 16MHz. As the peripherals continue to operate, pretty much any of them can wake the processor up.

"Stop" mode, on the other hand, has the potential to reduce power consumption to the sub-microamp level. This is the mode I used for the SmartProbe when it's 'off'. Here, the core and most peripheral clocks are halted; only the real-time clock and watchdog timer continue to run if they are enabled (which they aren't). Volatile memory is retained.

Several possible sources can wake the microcontroller from stop mode, including interrupts triggered by the states of external pins changing, which is what we use.

The final low-power mode is "Standby". In this mode, almost everything is powered off, including the RAM and almost all the registers. Only a very limited selection of wake-up sources is available, and the wake-up time is the longest.

Putting the processor into stop mode is in itself not enough to get the current consumption down to the very low levels we require. There is quite a lot to do both within the microcontroller and in the external circuit before executing the instruction that halts the processor.

Externally, we shut the OLED display off via an I^{2}C command, shut down the beeper PWM and turn off the boost converter. Internally, we stop the ADC conversion process and the timer that triggers it, disable the bandgap reference and the buffers that feed its output to the ADC and the comparator. We also limit the functionality of the accelerometer to just detecting activity or inactivity.

A close-up of the probe we used for our SmartProbe.

We also put most of the I/O pins into analog input mode so they look like high-impedance inputs, minimising any leakage currents that might otherwise occur. The data sheet suggests that the digital input schmitt trigger buffers can be a source of leakage – hence using the analog input mode.

In fact, leakage current is our number one enemy in ultra-low power circuits such as this, and it can come from some quite obscure sources.

For example, consider the 4.7kΩ I^{2}C pullup resistors. These are pulled up by one of the microcontroller's GPIO pins instead of being directly connected to the power rail (V_{BAT}) as one might do normally. Fig.6 shows why we have to do this. The I^{2}C bus connects to the micro and the accelerometer, which remain powered up in stop mode, but also to the display, which does not.

The display driver's I^{2}C inputs are internally protected by diodes connected as shown in the figure. In normal operation, these prevent the input pin rising more than about 0.6V above the 3.3V supply or falling more than 0.6V below ground.

However, when the boost converter is off, the display driver's positive power rail is at 0V, allowing a leakage current path from V_{BAT} to ground via the I^{2}C pullups, as shown by the red paths.

If V_{BAT} was at 2.5V and the protection diode forward drop was 0.5V, there would be more than 850µA leakage in total. This is 30 times more current than our target, so clearly not acceptable! Powering the pullups from a GPIO pin that can be put into a high-impedance state eliminates this problem.

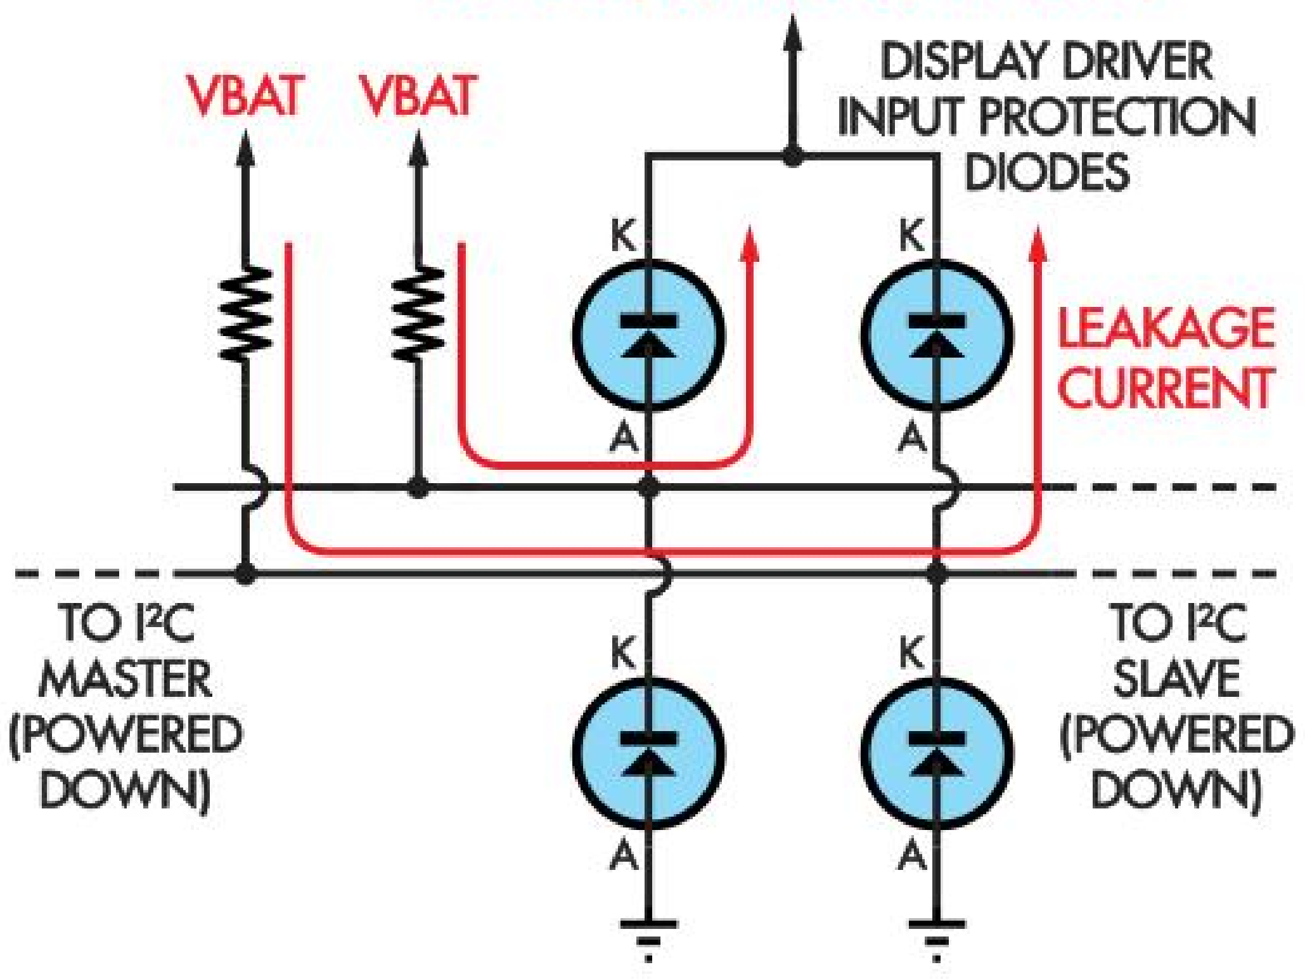
3.3V RAIL (POWERED DOWN TO 0V)
DISPLAY DRIVER INPUT PROTECTION DIODES
VBAT
VBAT
K
K
A
A
LEAKAGE CURRENT
TO I²C MASTER (POWERED DOWN)
K
K
TO I²C SLAVE (POWERED DOWN)
A
A

Fig.6: if the I^{2}C bus was pulled up to V_{BAT}, significant leakage currents would flow through the pullup resistors and display driver chip's input protection diodes when the 3V3 rail is off.

Once all the internal and external extraneous current consumers are dealt with, we are ready to stop the microcontroller core and peripherals. We just have to set or clear a few bits in various control registers to ensure the core enters the correct low-power mode and that it wakes up due to the right stimulus with the right clock source.

We have to disable all interrupt sources except the external interrupt pin associated with the accelerometer, and globally disable interrupts. Finally, we can execute a "Wait for Interrupt" (WFI) instruction that stops the core until an interrupt is received.

It might seem odd that we globally disable interrupts if we want to wake up due to an interrupt, but the way it works is that the peripheral's interrupt flag does the waking (in this case the external interrupt) regardless of the state of the global interrupt enable flag. By disabling global interrupts, we ensure that when the processor wakes up, it continues executing code where it left off and not in an ISR.

On resumption, we have to undo all the work we did before entering low-power mode, restoring the I/O pin states, enabling the boost converter and the internal regulators and buffers we disabled. We then reinitialise the drivers to start the ADC, the display and all the rest of the stuff necessary to resume operation.

Construction

The SmartProbe is built on a small double-sided printed circuit board coded P9054-04 that measures 54.5 × 29.5mm. Both sides of the board are fairly tightly packed with surface-mounting parts, although most are M2012/0805-sized (2.0 × 1.2mm), so pretty easy to handle. The overlay diagrams, Figs.7 & 8, show where everything goes.

Start assembly by mounting the trickiest part, the accelerometer (IC3). If you are using solder wire, you will need to apply some flux paste and a thin layer of solder to the pads first. Try to get roughly the same amount of solder on each pad – if there is too little on one or two, you risk an open-circuit connection. If you elect to use solder paste, try to get a more-or-less even smear across the footprint.

You can then reflow the chip using something like a hot air wand, making sure to get it the right way round. Despite its small size, I found this package fairly forgiving when it came to assembly.

Next, solder in the boost converter chip (REG4) and then the optional programming connector, CON3, if you will use it (it isn't necessary if you get a pre-programmed chip from our Online Shop). They should solder in fairly easily using a fine tip soldering iron, some good flux and a bit of solder wick to clean up any bridges.

Mount the rest of the components on the top side of the board, except the coin cell holder and the probe connector (CON1), in the order you prefer. I tend to install the finer-pitch and smaller parts first, working my way up to the larger ones.

Flip the board over and install the six passives on the back side, plus the audio transducer. Finally, add the coin cell holder and the probe connector (CON1) to the top side. It is a good idea to check your work and clean up the board with a bit of isopropyl alcohol or another solvent at this point, before installing the display.

The display's flat flex cable is soldered directly to the row of pads on the back of the board. Make sure the pin 1 designator on the flex is aligned with the small dot on the board. This should correspond with the display being face-up when folded back on itself. It should be face-down when the flat flex is straight, and it should extend over the side of the board nearest the row of pads.

Align the flat flex so that about 1mm of the PCB pads are visible, and secure it in place temporarily with a couple of bits of adhesive tape.

Double-check everything lines up and tack the display in place by soldering a couple of the pins. I suggest using two of the signal pins for this, rather than the pins connected to the ground plane, as they require a lot more heat.

If everything still looks good, go ahead and solder all the pins, taking care not to create any short circuits. You can remove the tape and clean up with a little solvent if you need to, but keep it away from the display itself.

If you intend to program the microcontroller in-circuit, do that now. You will need to make a suitable adaptor to connect the programmer to the flat flex connector. See the details in the accompanying panel.

You will have to insert a coin cell (or otherwise power the microcontroller) while it is being programmed. If your microcontroller came preprogrammed, you should skip this step.

You can now perform a quick test by inserting a coin cell into the holder. You should hear the start-up beep, and the display should come to life in voltage mode with a reading close to 0.000V. The least significant digit (millivolts) and the sign may move around a little, but the rest of the digits should be zero.

If you do nothing, the probe should

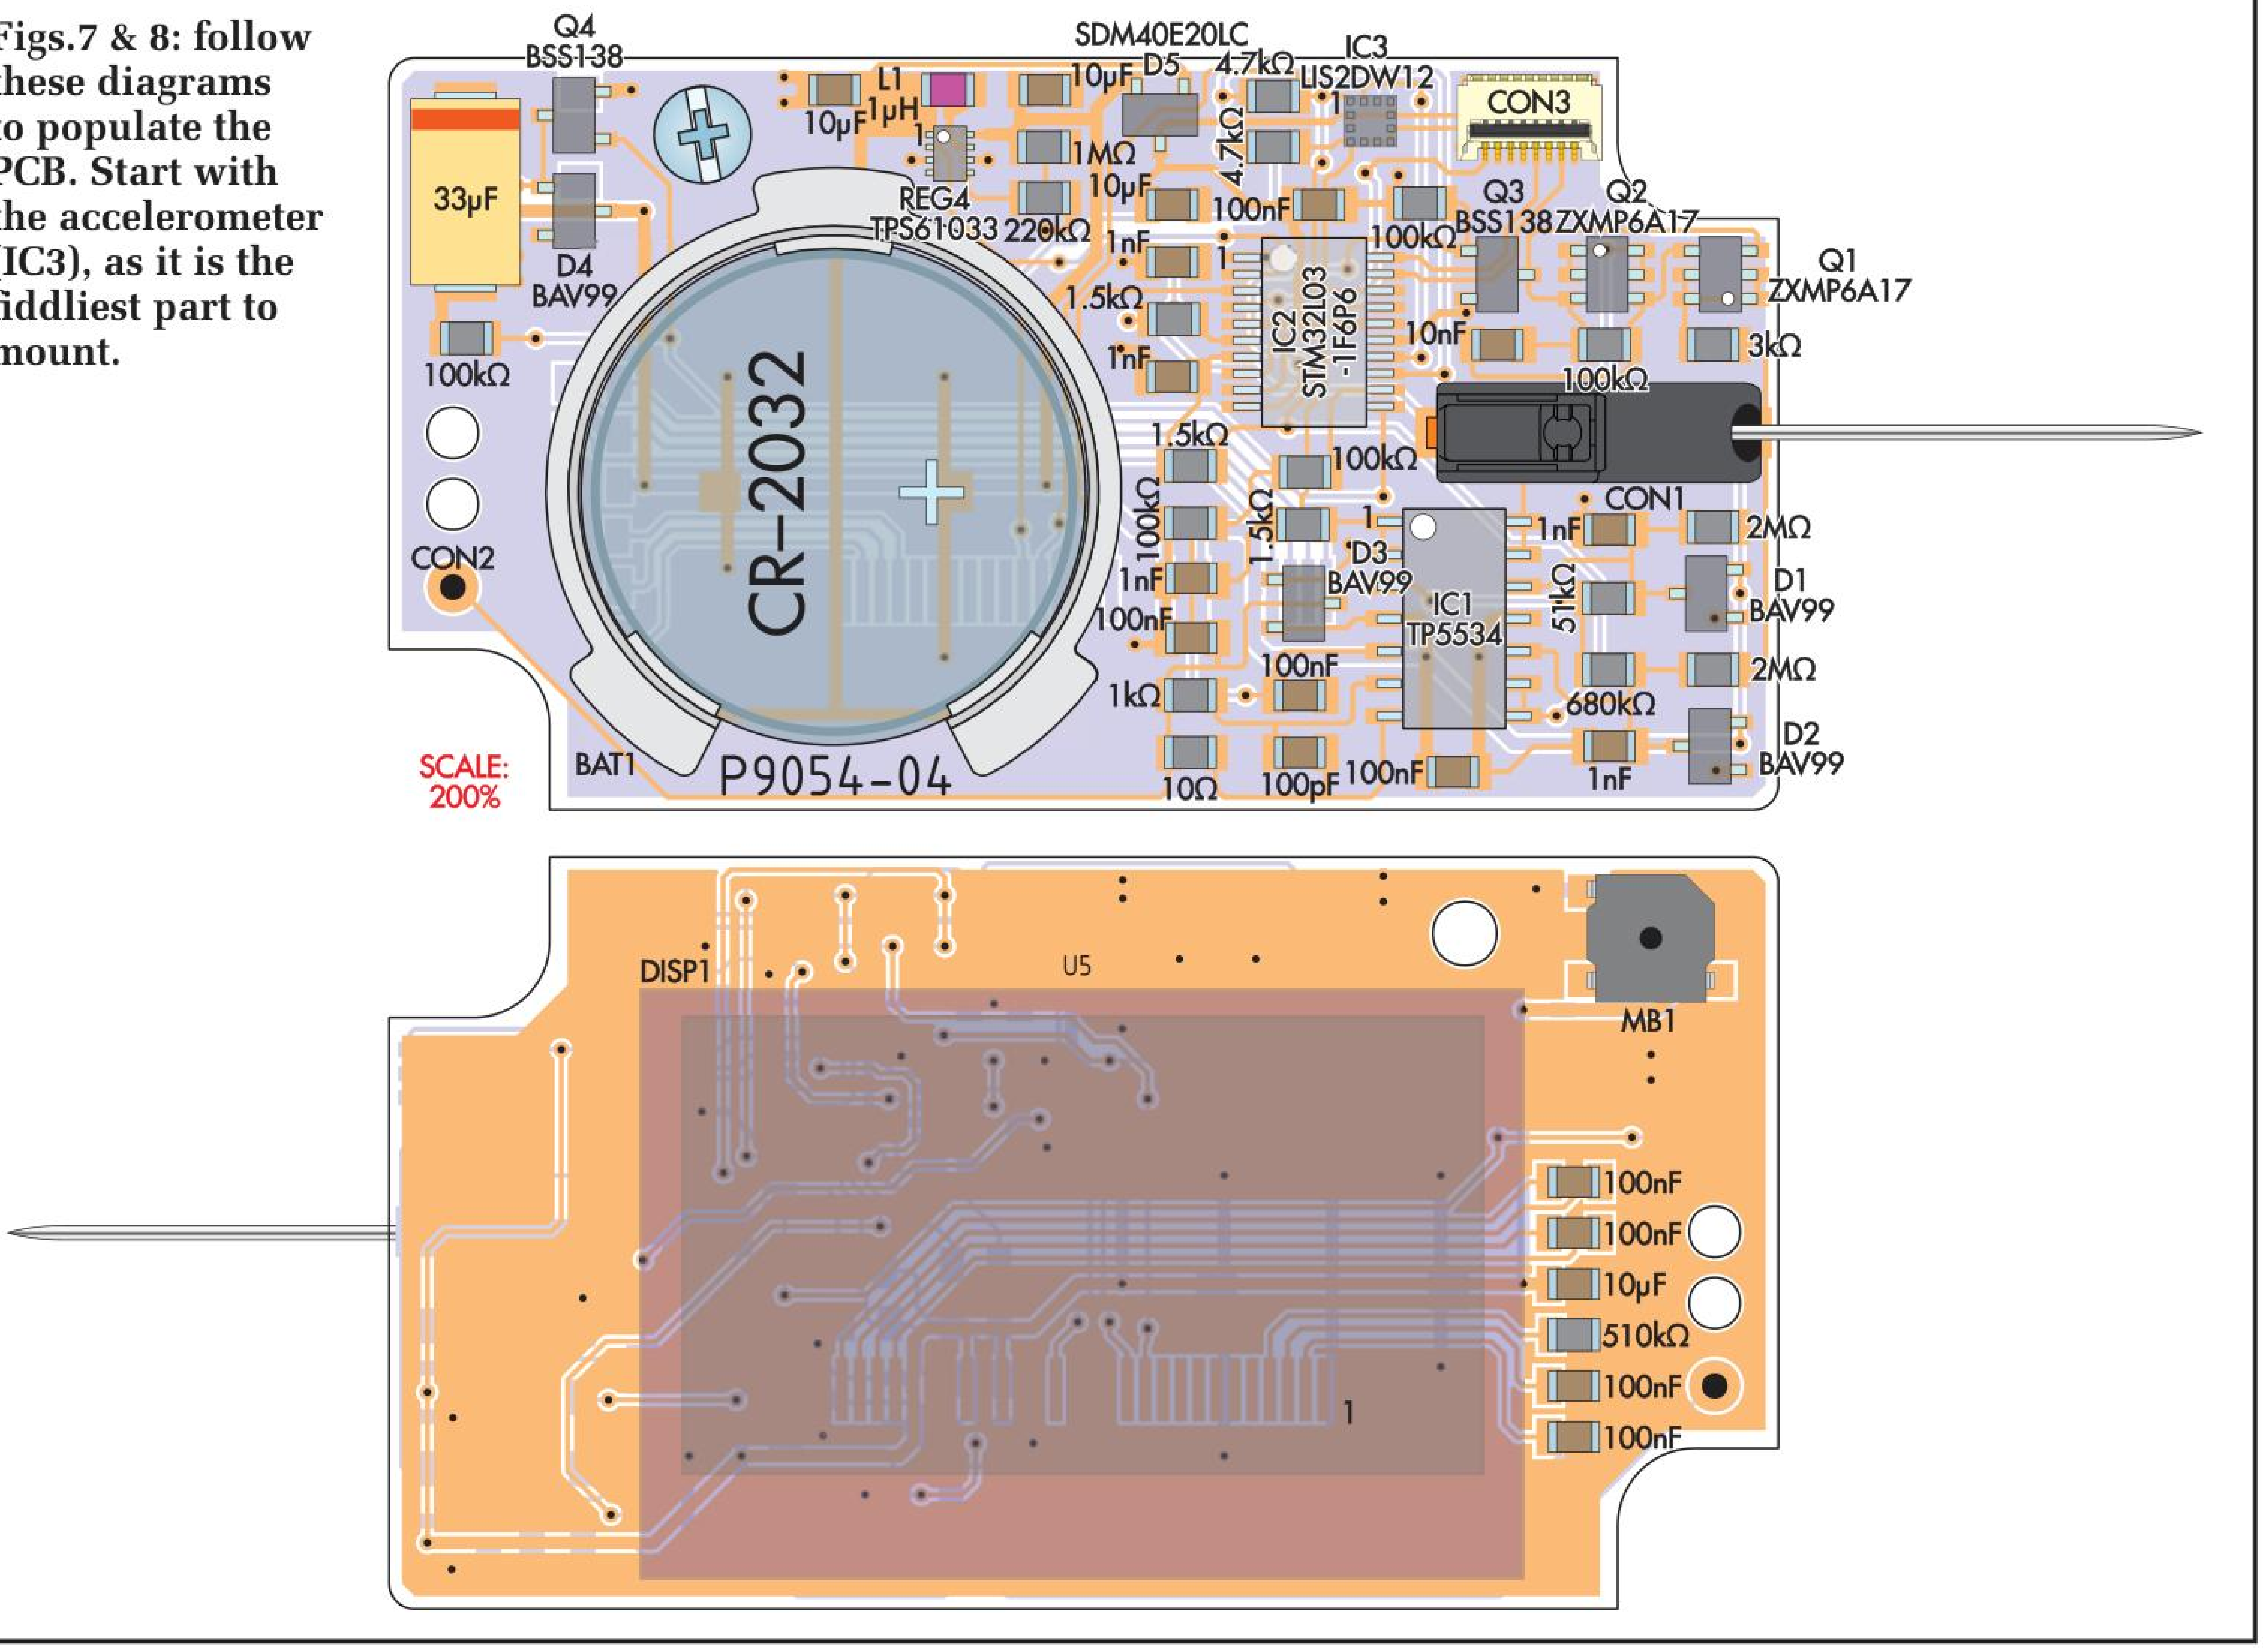
Q4
BSS138
33µF
D4
BAV99
100kΩ
CON2
L1
1µH
10µF
REG4
TPS61033
220kΩ
1MΩ
10µF
SDM40E20LC
D5
4.7kΩ
4.7kΩ
IC3
LIS2DW12
CON3
100nF
100kΩ
1nF
1.5kΩ
1nF
IC2
STM32L03
1F6P6
Q3
BSS138
Q2
ZXMP6A17
Q1
ZXMP6A17
10nF
3kΩ
100kΩ
CR-2032
1.5kΩ
100kΩ
100kΩ
1.5kΩ
1nF
100nF
D3
BAV99
IC1
TP5534
CON1
1nF
51kΩ
2MΩ
D1
BAV99
2MΩ
100nF
1kΩ
680kΩ
D2
BAV99
SCALE: 200%
BAT1
P9054-04
10Ω
100pF
100nF
1nF
DISP1
U5
MB1
100nF
100nF
10µF
510kΩ
100nF
100nF
1

Figs.7 & 8: follow these diagrams to populate the PCB. Start with the accelerometer (IC3), as it is the fiddliest part to mount.

switch itself off after five or six seconds, and it should wake up again you give it a bit of a jiggle.

This is enough to test the basic functionality. If everything is OK, you can remove the coin cell and fix the display in place with a small piece of double-sided foam tape. Fold it over and align the edges of the display glass with the outer rectangle on the PCB silkscreen.

Debugging

If there is a problem, do some troubleshooting before fixing the display in place. If the unit appears dead, first measure the 3.3V rail to make sure it is working when the unit is awake.

If 3.3V is not present, check that the battery voltage is getting to it and the boost converter enable pin is high. The latter is a sign that the micro is at least trying to start the converter and means the problem is in the boost converter itself or there is a short on the 3.3V rail.

If the 3.3V rail is fine and you hear the start-up beep, but the screen remains blank, the problem is probably in the I^{2}C bus or with the display. Check the associated components and the soldering of the display connector. It is possible to get shorts under the flat flex that you can't necessarily see. Use a multimeter to look for (unwanted) shorts between adjacent pins if you suspect this may be the case.

The firmware does have a fair bit of error-detection built in. If an error occurs, the beeper emits a short, low tone (it is fairly quiet, so listen carefully) and displays a small fault icon on the screen.

If this says "ACC!", the code encountered a problem communicating with the accelerometer. If it says "DIS!", the problem is with the display communication, and if it says "SYS!" the problem is a processor exception, so probably related to corrupt code.

Mechanical assembly

Once everything is working as expected, you can prepare the case. Mark out and cut the opening for the display and the sound hole according to Fig.9. The display aperture can be cut by drilling a series of holes inside the marking and finishing up to the line with files. I used a file to put a chamfer around the display hole – the exact dimensions are not critical as this is purely cosmetic.

Drill the holes in each end of the

Parts List – SmartProbe

1 Hammond 1551JBK 60 × 35 × 15mm black ABS enclosure
[Farnell 1877310, RS 7417398, TME]
1 double-sided PCB coded P9054-04, 54.5 × 29.5mm
1 BC-2600 SMD CR2032 cell holder (BAT1)
1 CR2032 3V lithium coin cell (BAT1)
1 Wago 2946-2060-471/998-404 single entry SMD terminal block (CON1)
1 Molex 503480-0800 8-pin flat flex connector, 0.5mm pitch (CON3; **optional**, for ICSP)
1 X30654 128×64-pixel 3.3V graphic OLED screen with SH1106 controller (DISP1) **[AliExpress 32890183042]**
1 WLPH201610M1R0PP or equivalent 1μH 2A+ inductor, SMD M2012/0805 size (L1)
1 CMT-0525-75-SMT-TR SMD magnetic transducer buzzer (MB1)
1 6mm-long No.6 self-tapping screw
1 size 4 straw sewing needle
1 alligator clip
1 length of medium-duty hookup wire
1 short length of 2mm diameter heatshrink tubing
1 small piece of double-sided foam-cored tape

Semiconductors
1 TP5534-SR quad chopper-stabilised op amp, SOIC-14 (IC1)
1 32-bit STM32L031F6P6 microcontroller programmed with 0411025A.HEX, TSSOP-20 (IC2)
1 LIS2DW12 ultra-low-power 3-axis accelerometer, LGA-12 2 × 2mm (IC3)
1 TPS61033DRLR adjustable boost regulator, SOT-583 (REG4)
2 ZXMP6A17E6TA 60V 2.7A logic-level P-channel Mosfets, SOT-23-6 (Q1, Q2)
2 BSS138K 50V 220mA N-channel logic-level Mosfets, SOT-23 (Q3, Q4)
4 BAV99 ultrafast dual series diodes, SOT-23 (D1-D4)
1 SDM40E20LC 20V 400mA dual common-cathode schottky diode, SOT-23 (D5)

Capacitors *(all SMD M2012/0805 size 50V X7R ceramic unless noted)*
1 33μF 16V D-case tantalum
4 10μF 16V
8 100nF
1 10nF
5 1nF
1 100pF C0G/NP0

Resistors *(all SMD M2012/0805 size, 1% unless noted)*
2 2MΩ ±0.1%
1 1MΩ
1 680kΩ ±0.1%
1 510kΩ
1 220kΩ
5 100kΩ
1 51kΩ ±0.1%
2 4.7kΩ
1 3kΩ
3 1.5kΩ
1 1kΩ
1 10Ω

Optional SWD flat flex adaptor

1 double-sided PCB coded P9045-A, 35 × 25mm
1 Molex 503480-0800 8-pin flat flex connector, 0.5mm pitch (CON1)
1 CNC Tech 3220-10-0300-00 10-pin, 1.27mm-pitch box header (CON2)
1 Molex 0150200079 8-way, 0.5mm pitch 100mm flat flex cable

SWD Programming Adaptor

There was no space on the main PCB for the standard ST Micro SWD/JTAG programming header, which is a 2×5-pin miniature shrouded box header (1.27mm pin pitch). Thus, a more compact 8-way FFC connector was used. This small PCB adapts that to the more standard connector so that a programmer can be connected via a ribbon cable with an IDC plug.

Its circuit is shown in Fig.a. Because pins 7 & 8 are not connected to the main board, only the single-wire debug (SWD) protocol is supported, not JTAG. Importantly, note that the pinout of CON1 is reversed compared to CON3 in Fig.2. That's because the FFC is inserted flat, such that the connections are reversed between those two connectors.

The adaptor is built on a double-sided PCB that's coded P9045-A and measures 35 × 25mm. The assembly should be straightforward, referring to the overlay diagram, Fig.b. The TP-X pin is provided to allow a custom debugging signal to be generated from the microcontroller.

When using it, make sure that the FFC cable is inserted with the correct orientation between the two ends. The easiest way to check this is to verify continuity between the grounds of the SWD Adaptor board and target board before connecting the programmer.

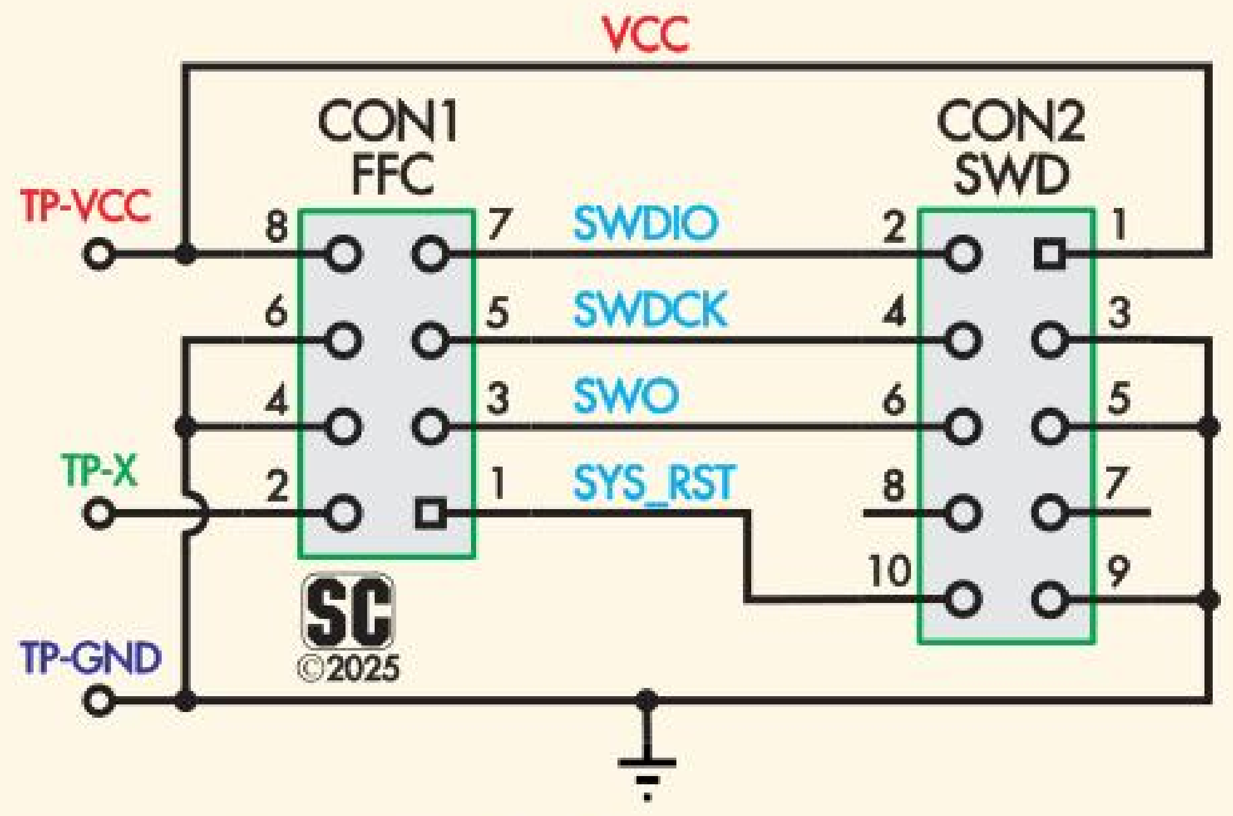

Fig.a: this SWD adaptor connects to the main circuit (Fig.2) via a flat flex cable (FFC) and allows a standard ST Micro programmer to connect via CON2. The FFC is connected such that it reverses the connections, making CON1's pinout correspond to CON3 in Fig.2.

A close-up of the Adaptor's flex cable is shown above and the finished PCB below. We have used some tape on both ends to provide extra rigidity.

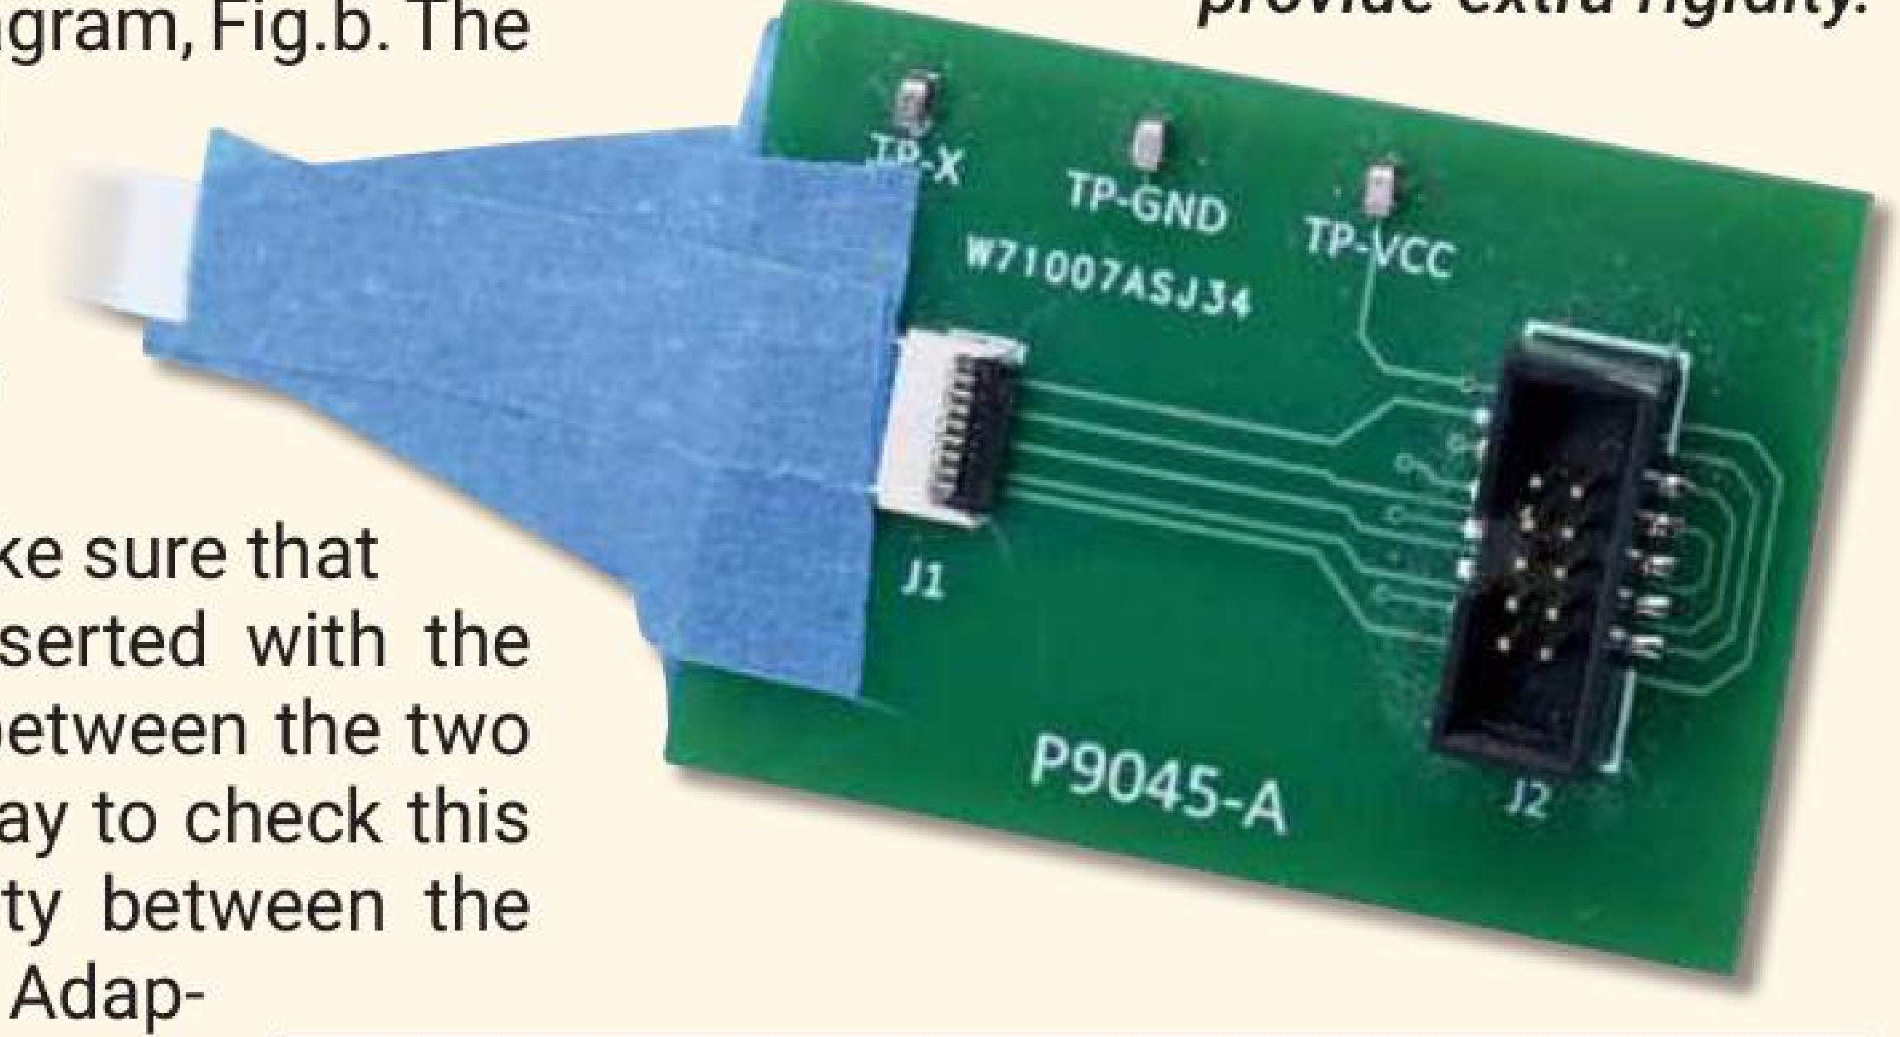

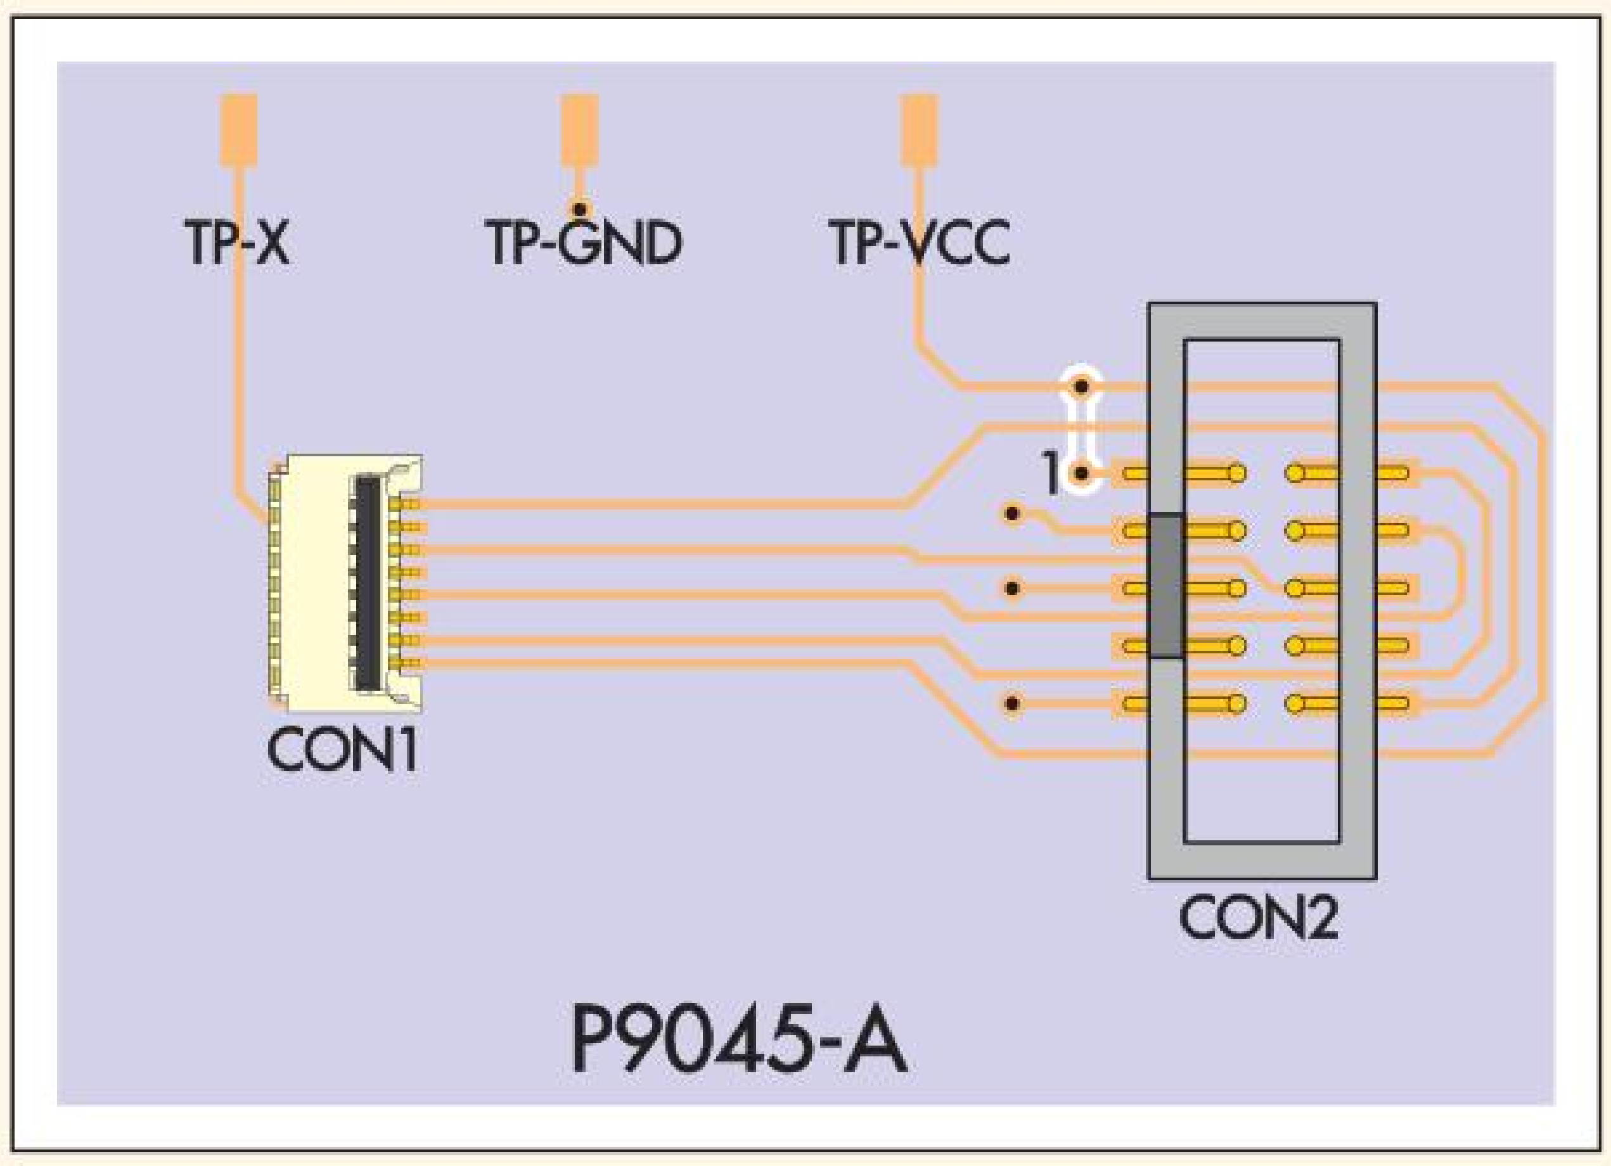

Fig.b: the SWD adaptor PCB has just two components, the connectors, plus three test points.

case to suit the diameter of the ground lead and probe you choose (the dimensions shown are what we used). Try to make the probe hole a close, but not tight, fit with the probe, including its insulating sleeve. If it is too loose, the probe may wobble around.

I made my probe from a sewing needle. These are nice and sharp, so good for probing surface-mount parts, and fairly hard, so they last a while. I used a size 4 straw-type needle, but any needle with a diameter between 0.5mm and 1.0mm should do. I covered the needle in heatshrink tubing, leaving it about 5mm short of either end.

The eye of my needle was a similar diameter to the shaft, so I could insert it through the case and into the connector as-is. You may have to cut the eye off your needle if it is wider than about 1mm; otherwise, it may be too big for the connector.

The connector I used for the probe has a spring operation. You press down on the small divot at the top to insert or remove the conductor. When released, the connector firmly grips the probe. That, plus the close fit of the hole in the case, is enough to hold the probe solidly in place.

Finally, you need to clip out one of the two PCB mount bosses inside the case as it interferes with the display. The one to remove is diagonally opposite the sound hole.

The PCB is secured to the case with a #6 × 6mm self-tapping screw – see Figs.10 & 11. You should solder the ground lead to the PCB before you finally mount the board. It solders in from the top (battery) side, then loops down and up again through the two strain relief holes to emerge on the same side. Thread the wire through the hole in the case before installing the board into the case.

I used a ground lead with a connector on the far end, so I had to thread the near end through the case before soldering it.

Remember to secure the lid to the case with the two screws provided. Coin cells are extremely dangerous for children, and it is mandatory that they are only accessible with the use of a tool. You should also be very careful about where you store them. Keep them in their special child-resistant packaging until they are required and always well and truly away from inquisitive little hands.

Using the SmartProbe

There is no need to calibrate the SmartProbe, but you can check its operation fairly easily with basic test equipment. The absolute accuracy is measured by setting a bench supply to some voltage near the middle of each range (I used 2V and 30V) and comparing the SmartProbe reading to that from a known good meter, preferably one with five or more digits.

The three units I built were all well within ±25mV on the 5V scale and 250mV on the 50V scale (0.5% of full scale in each case).

You can get an idea of its precision by taking a series of readings (the more the better) and looking at the variation between them. Successive readings should not differ more than about

±5mV and ±50mV on the two ranges, respectively. Keep the multimeter connected when you do this, to make sure the measured voltage does not change.

I measured the current consumption of the three units I built. The maximum current when on was between 15mA and 25mA, depending on the cell voltage. This means the average battery life when on will be around 12 hours. When off, the current consumption was always less than 25µA, corresponding to a shelf life of about 387 days.

Of course, neither of these is a realistic scenario. However, if we assumed an on-time of 30 minutes per week (remember it switches off after five or six seconds of inaction, so this is 30 minutes of actual measurement time), a fresh cell should last a little over 4 months.

The SmartProbe won't replace my multimeters, but if it becomes the first piece of test equipment that I reach for, I will consider that a success! **PE**

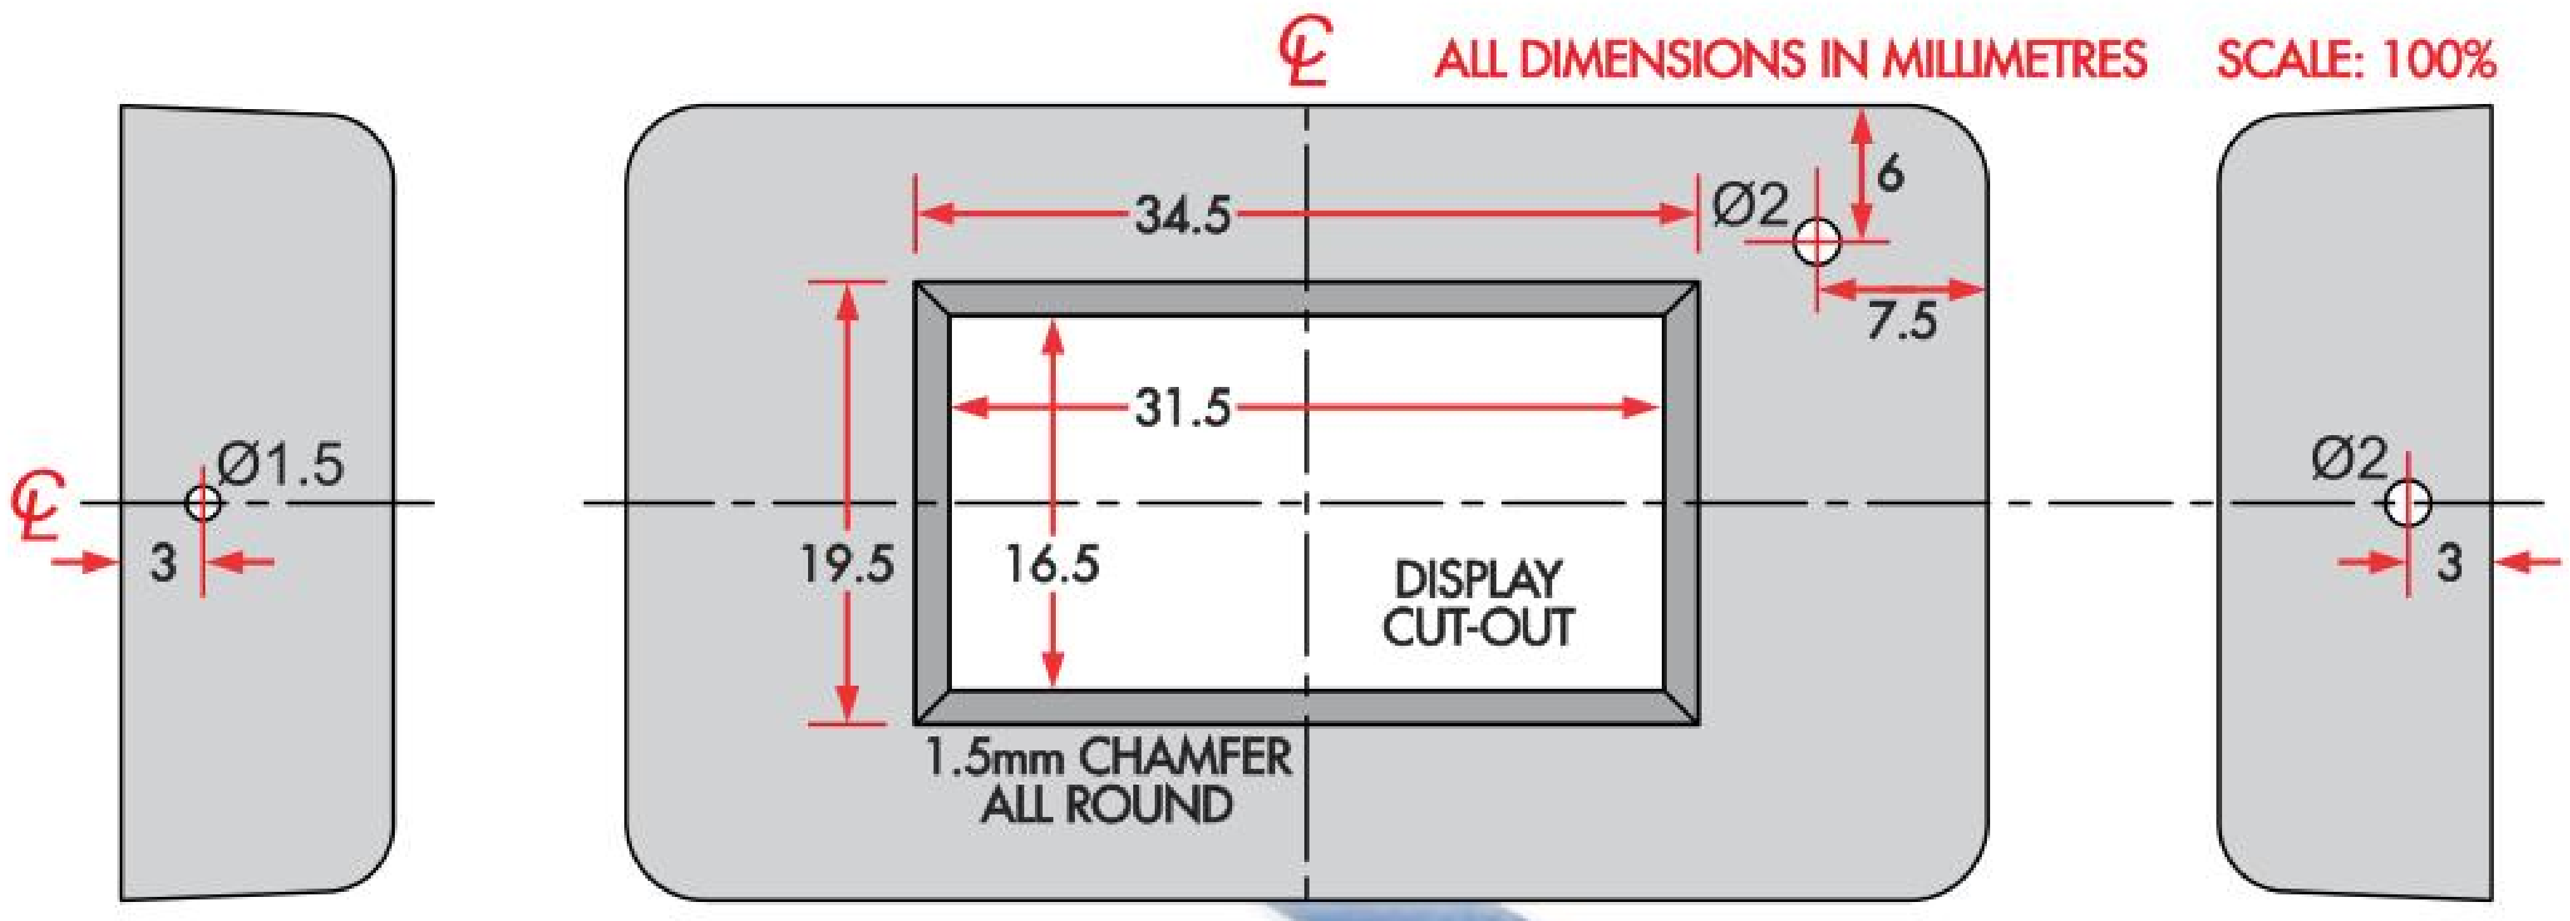
℄
ALL DIMENSIONS IN MILLIMETRES
SCALE: 100%
34.5
Ø2
6
7.5
31.5
Ø1.5
℄
3
19.5
16.5
DISPLAY CUT-OUT
1.5mm CHAMFER ALL ROUND
Ø2
3

Fig.9: drill the holes in the ends of the case to sizes that suit your probe and ground wire. You want a close (but not tight) fit, so the probe is held firmly in place.

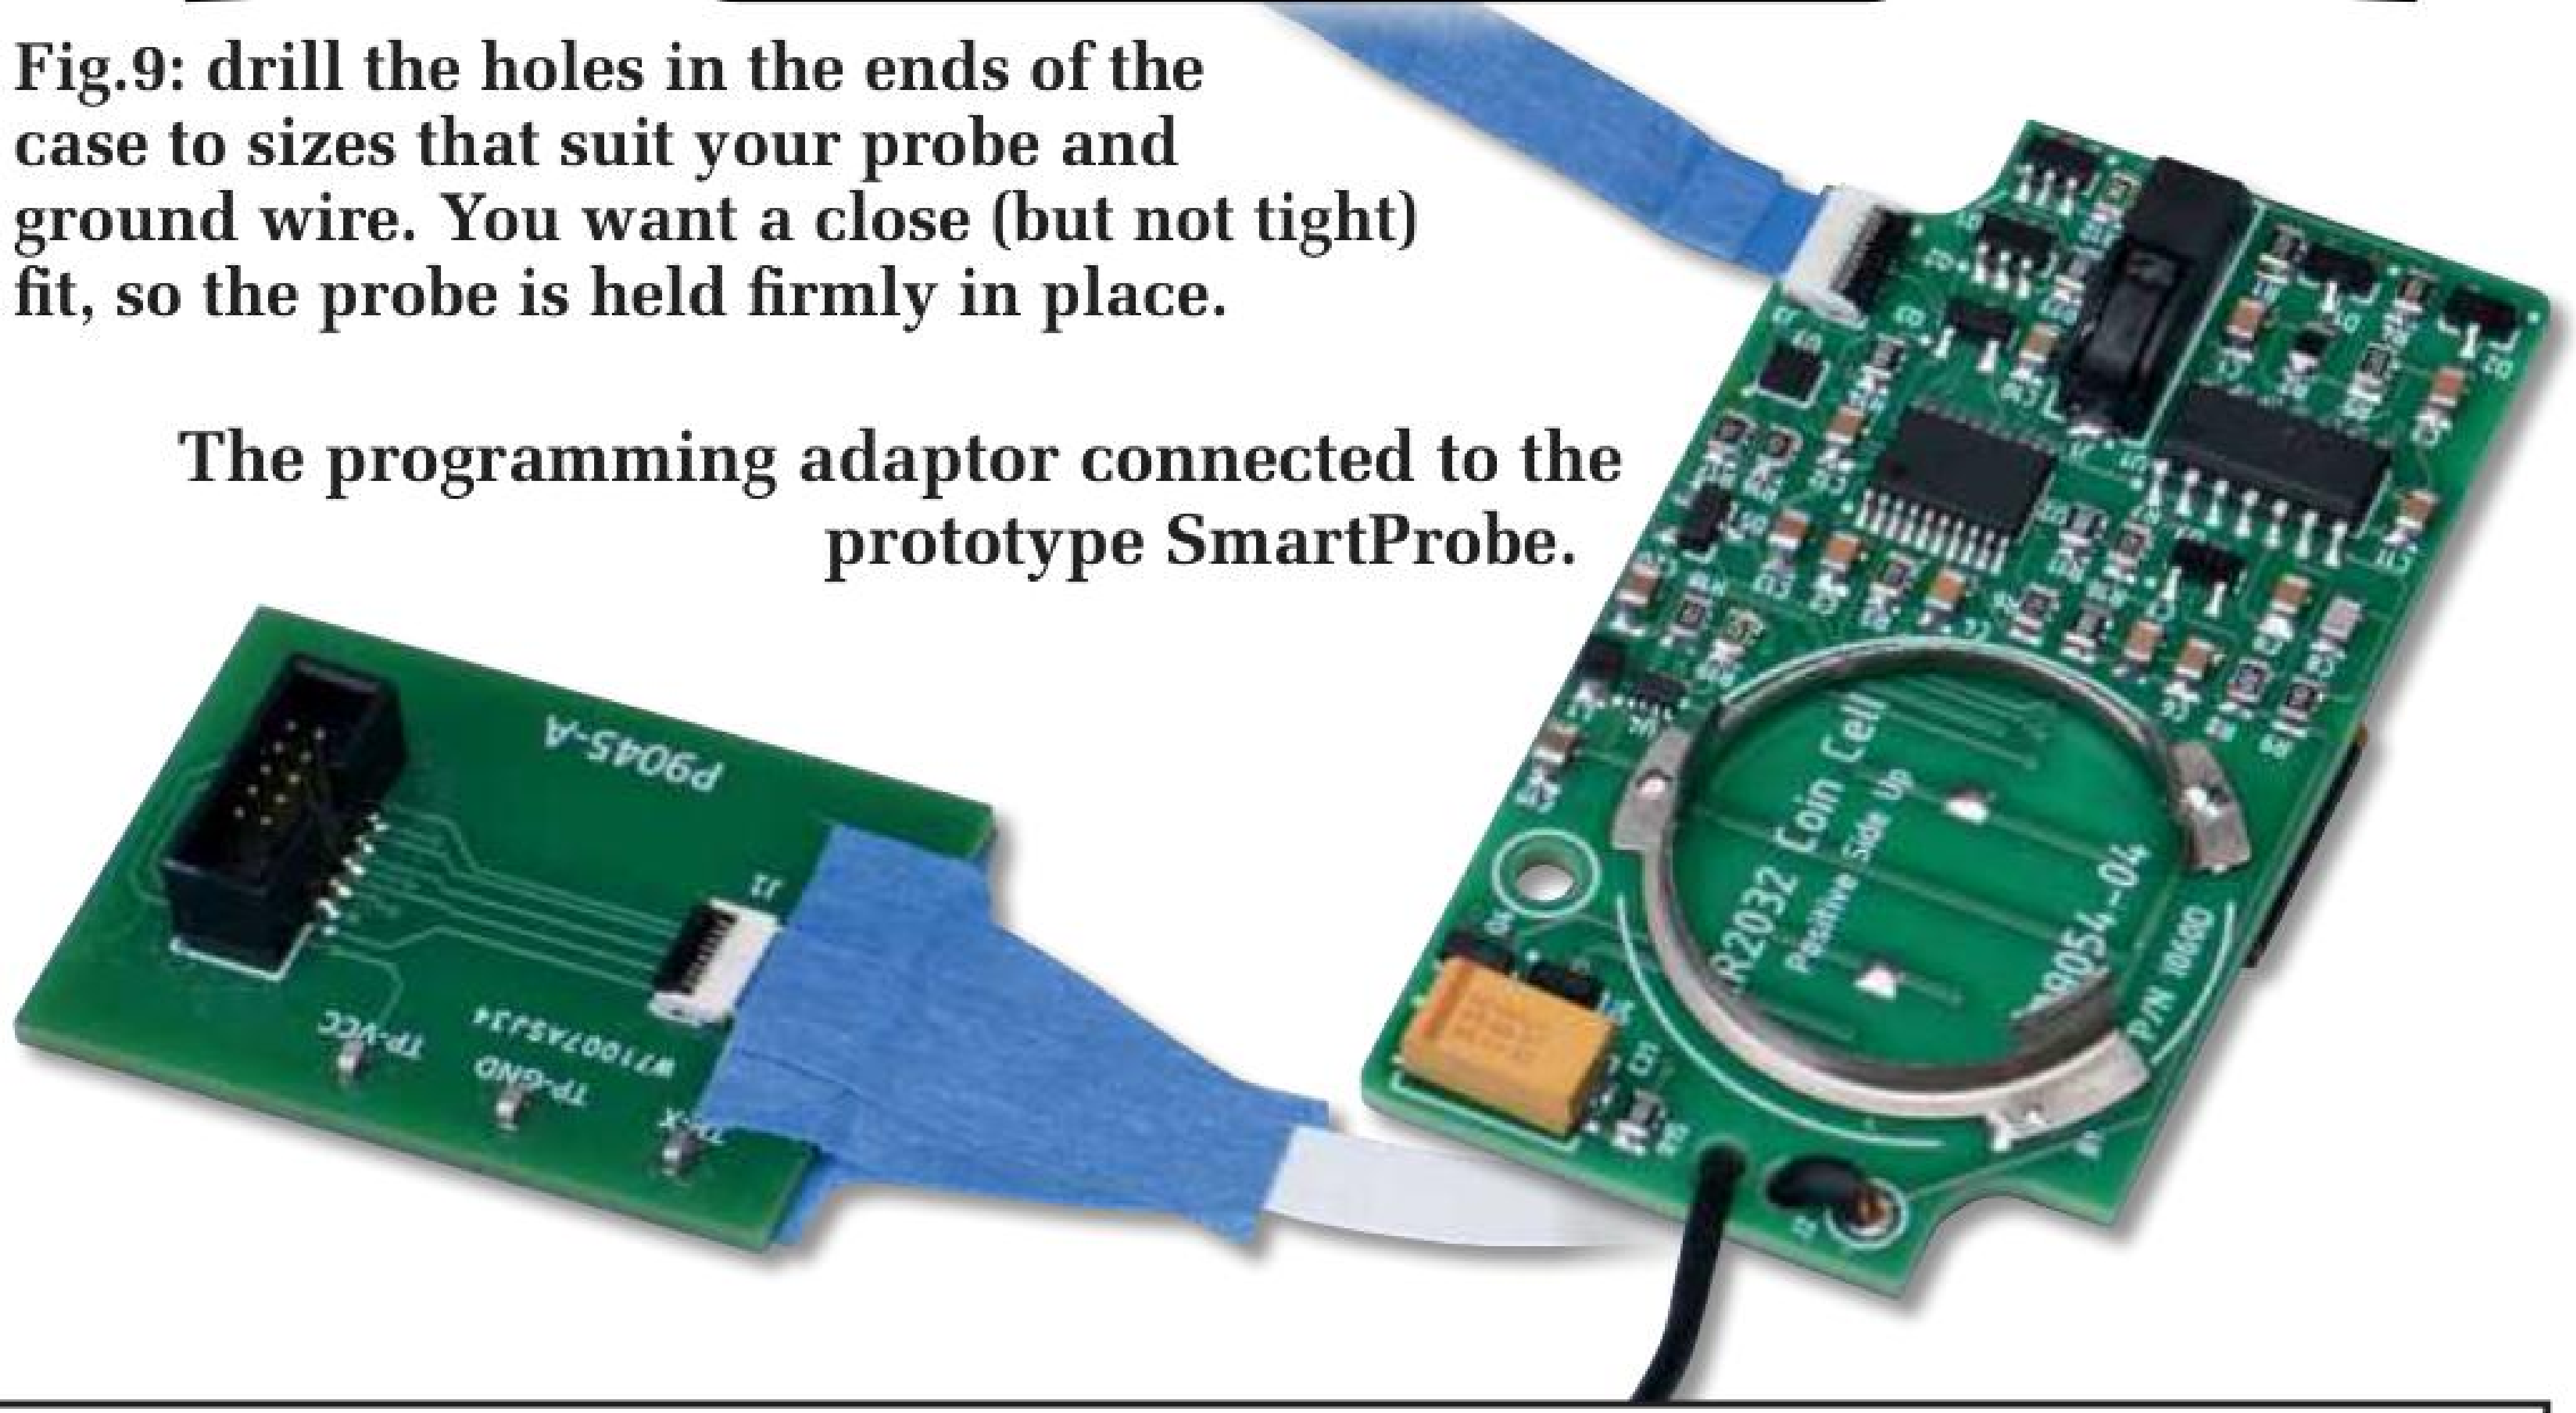
P9045-A
TP-VCC
TP-GND
R2032 Coin Cell
Positive Side Up
P9054-04

The programming adaptor connected to the prototype SmartProbe.

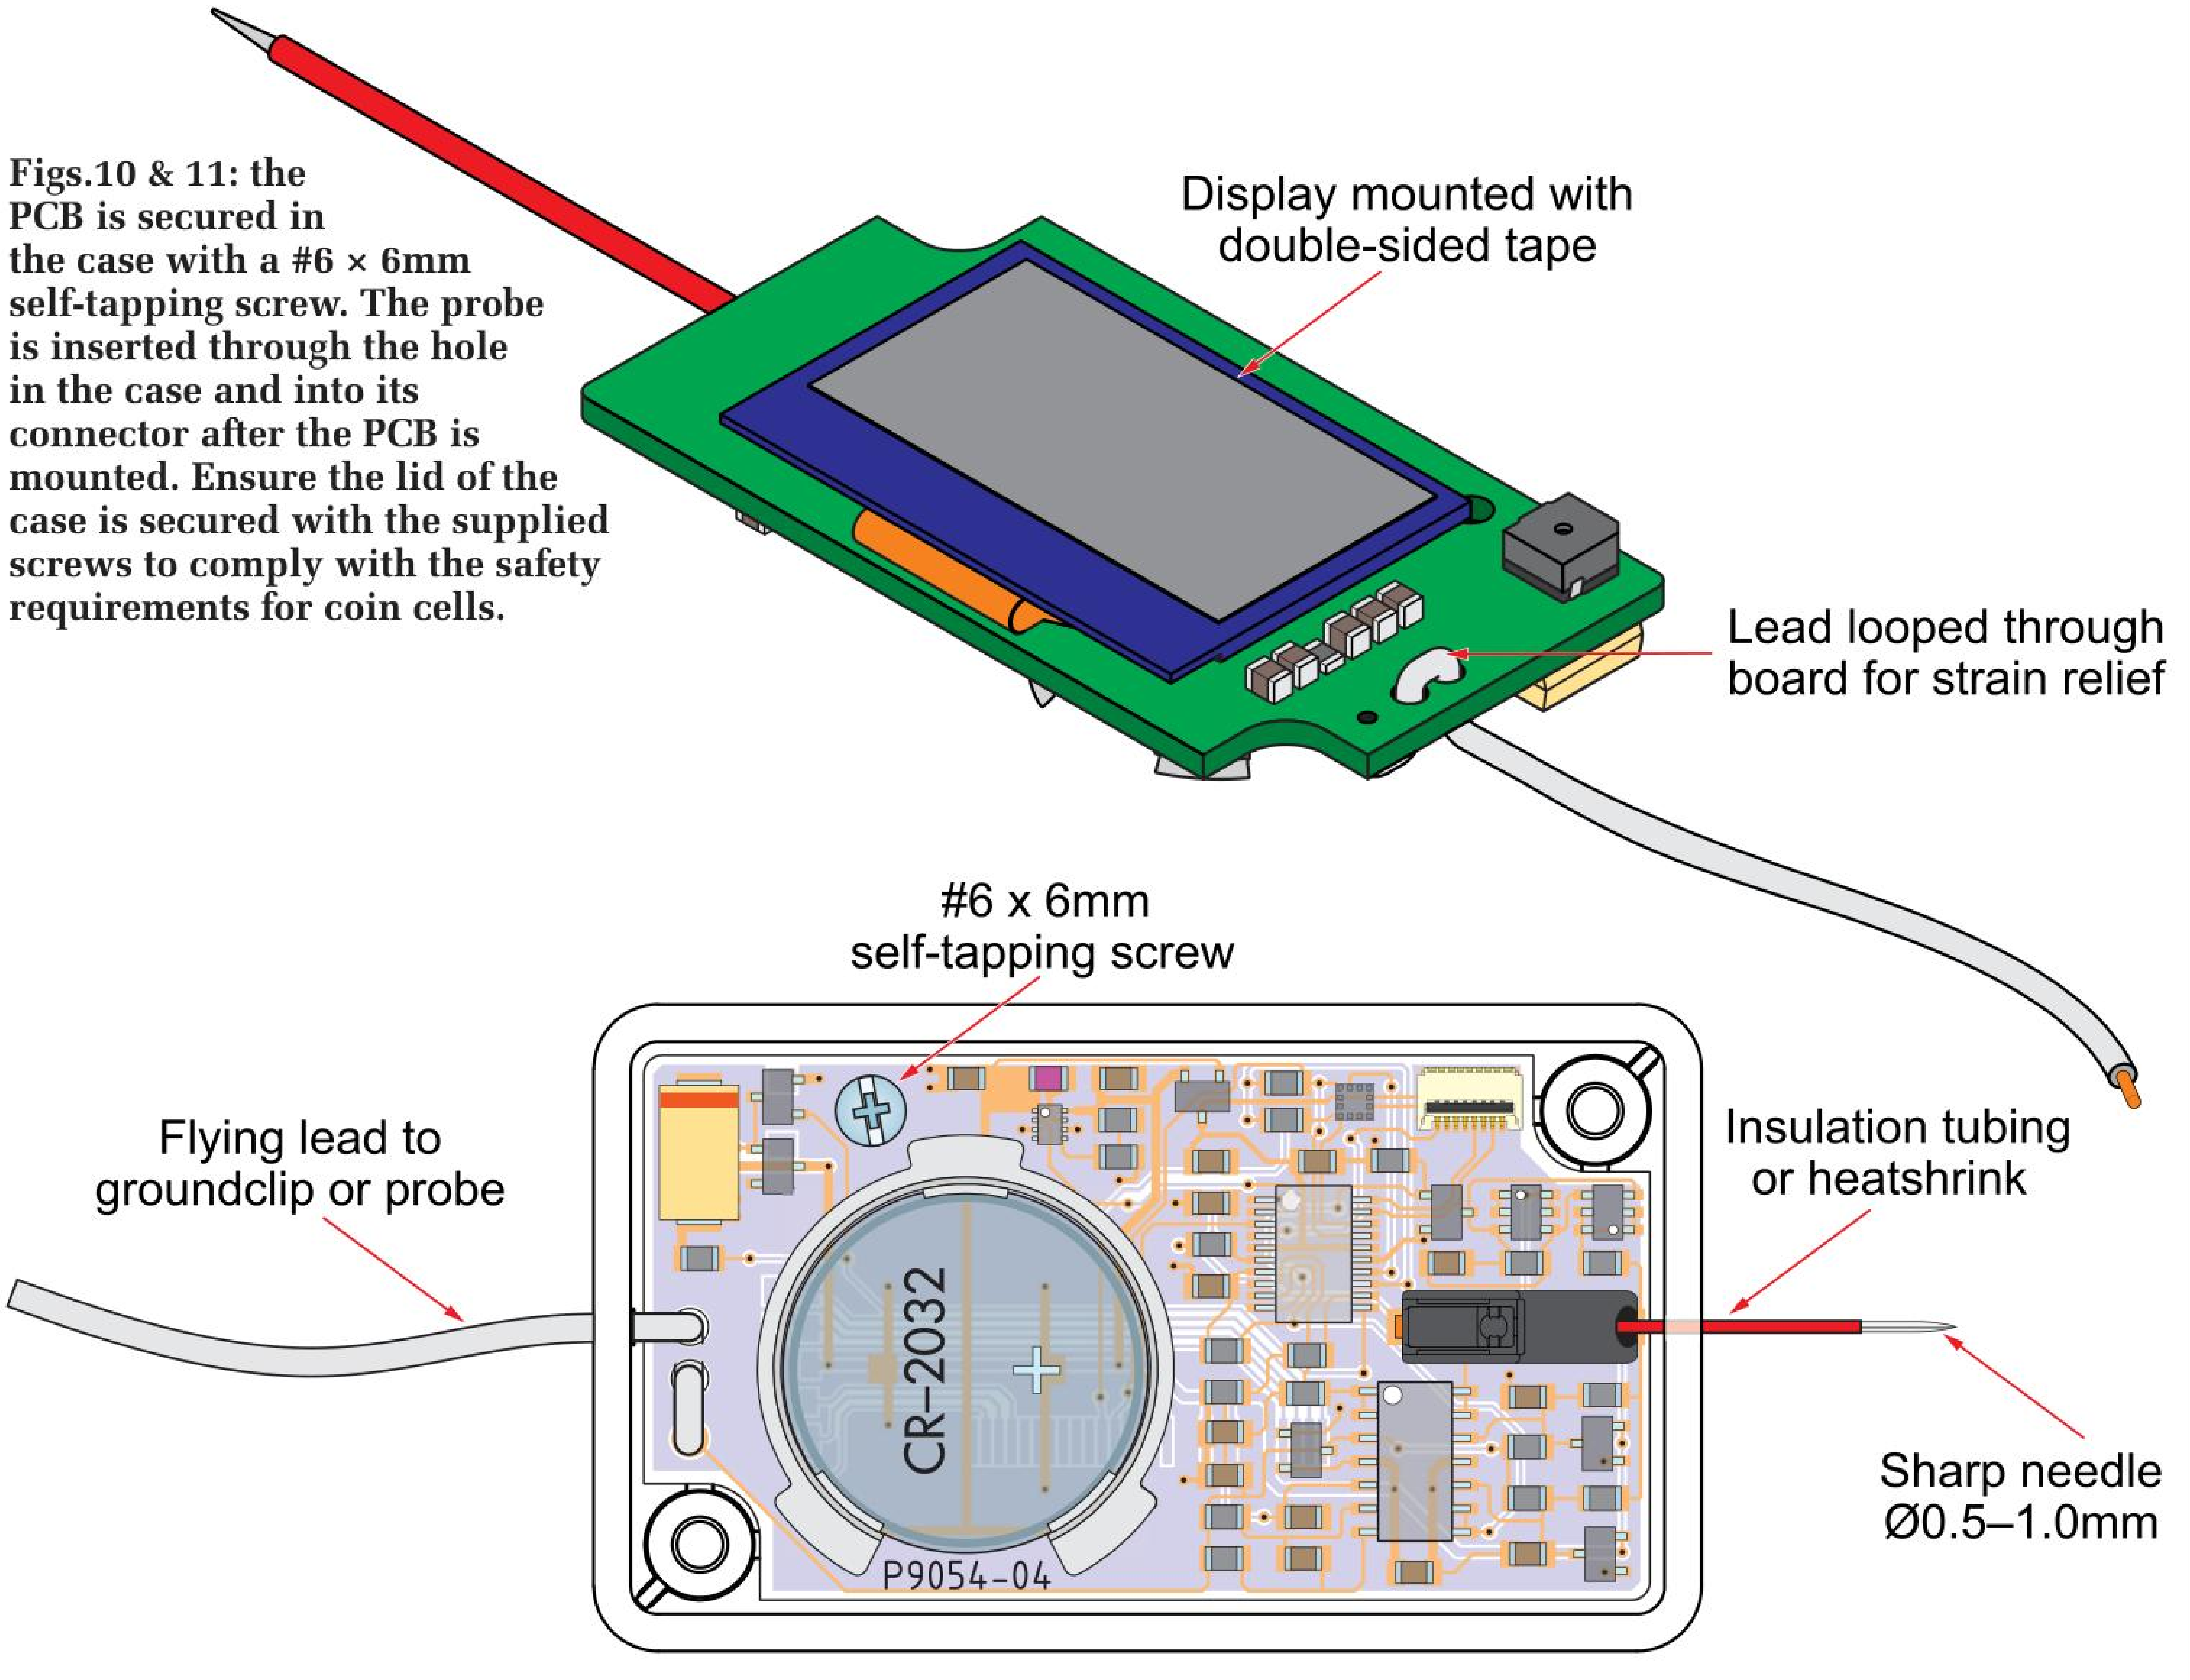
Display mounted with double-sided tape
Lead looped through board for strain relief
#6 x 6mm self-tapping screw
Flying lead to groundclip or probe
Insulation tubing or heatshrink
Sharp needle Ø0.5–1.0mm
CR-2032
P9054-04

Figs.10 & 11: the PCB is secured in the case with a #6 × 6mm self-tapping screw. The probe is inserted through the hole in the case and into its connector after the PCB is mounted. Ensure the lid of the case is secured with the supplied screws to comply with the safety requirements for coin cells.

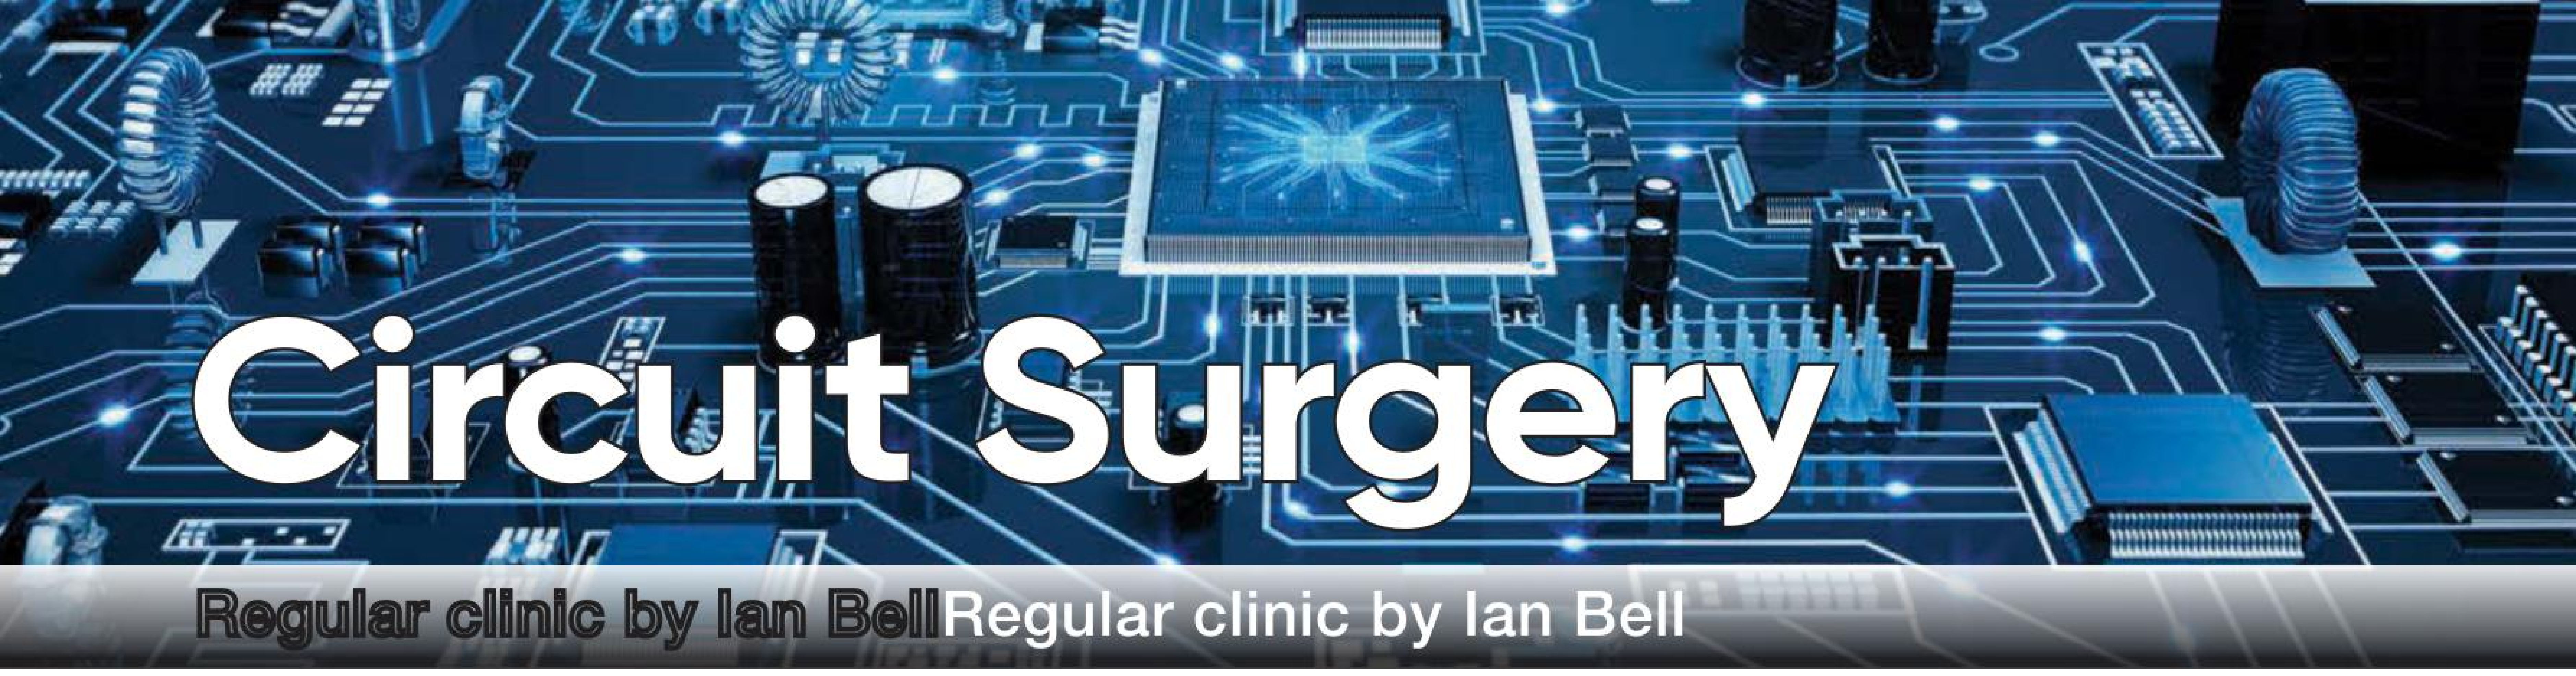

Circuit Surgery

Regular clinic by Ian Bell

Regular clinic by Ian Bell

Measuring circuit frequency response using a PC sound card, part 10: differential amplifiers and line drivers

In this series, we have been looking at circuit frequency response, including a measurement system using the 'sound card' (audio interface) of a PC or laptop together with free software called REW (Room EQ Wizard).

This was motivated by the need to measure the response of a practical implementation of an example digital filter (from our series on DSP) without using advanced lab test equipment.

Over the last few months, we have described parts of the signal conditioning circuitry that might be needed to interface the circuit under test to the sound card's line inputs and outputs. Last month, we started looking at the THAT1646 balanced line driver, which produces the output to the sound card in our circuit.

We introduced that device previously but mainly focused on the basic principles of balanced lines and differential signals. We also looked at some problems with simulating them.

This month, we will go over more details of the THAT1646 and discuss differential amplifiers in general. As with our previous discussions, these topics have wider applicability than just the frequency response measurement system.

The THAT1646 audio differential line driver is available in DIP-8, SO-8 and SO-16W (SO-16 comes in narrow and wide [W] variants) packages.

Fig.1 shows the internal circuit structure and pinout for the DIP-8 and SO-8 versions. The resistors and capacitors form the feedback network, which is significantly more complex than a basic op amp amplifier due to there being two outputs and the need for both common-mode and differential feedback.

The SNS± inputs can be directly connected to their respective outputs, but connecting them via capacitors reduces the common-mode output (as described in the datasheet).

The THAT1646's internal circuitry uses a fully differential amplifier, with differential inputs and outputs. However, one of its 5kΩ input resistors is internally grounded, so the input to the THAT1646 is single-ended. A similar device, the THAT1606, has both differential inputs and outputs.

THAT1646 simulation

Fig.2 shows an LTspice simulation for a minimal THAT1646 circuit that we started to investigate last month. The 100kΩ load resistors (**R1** and **R2**) and input buffer amplifier (**U1**) provide more realistic conditions than directly driving the device from an ideal source and leaving the outputs unloaded. An input buffer would be needed in many real designs,

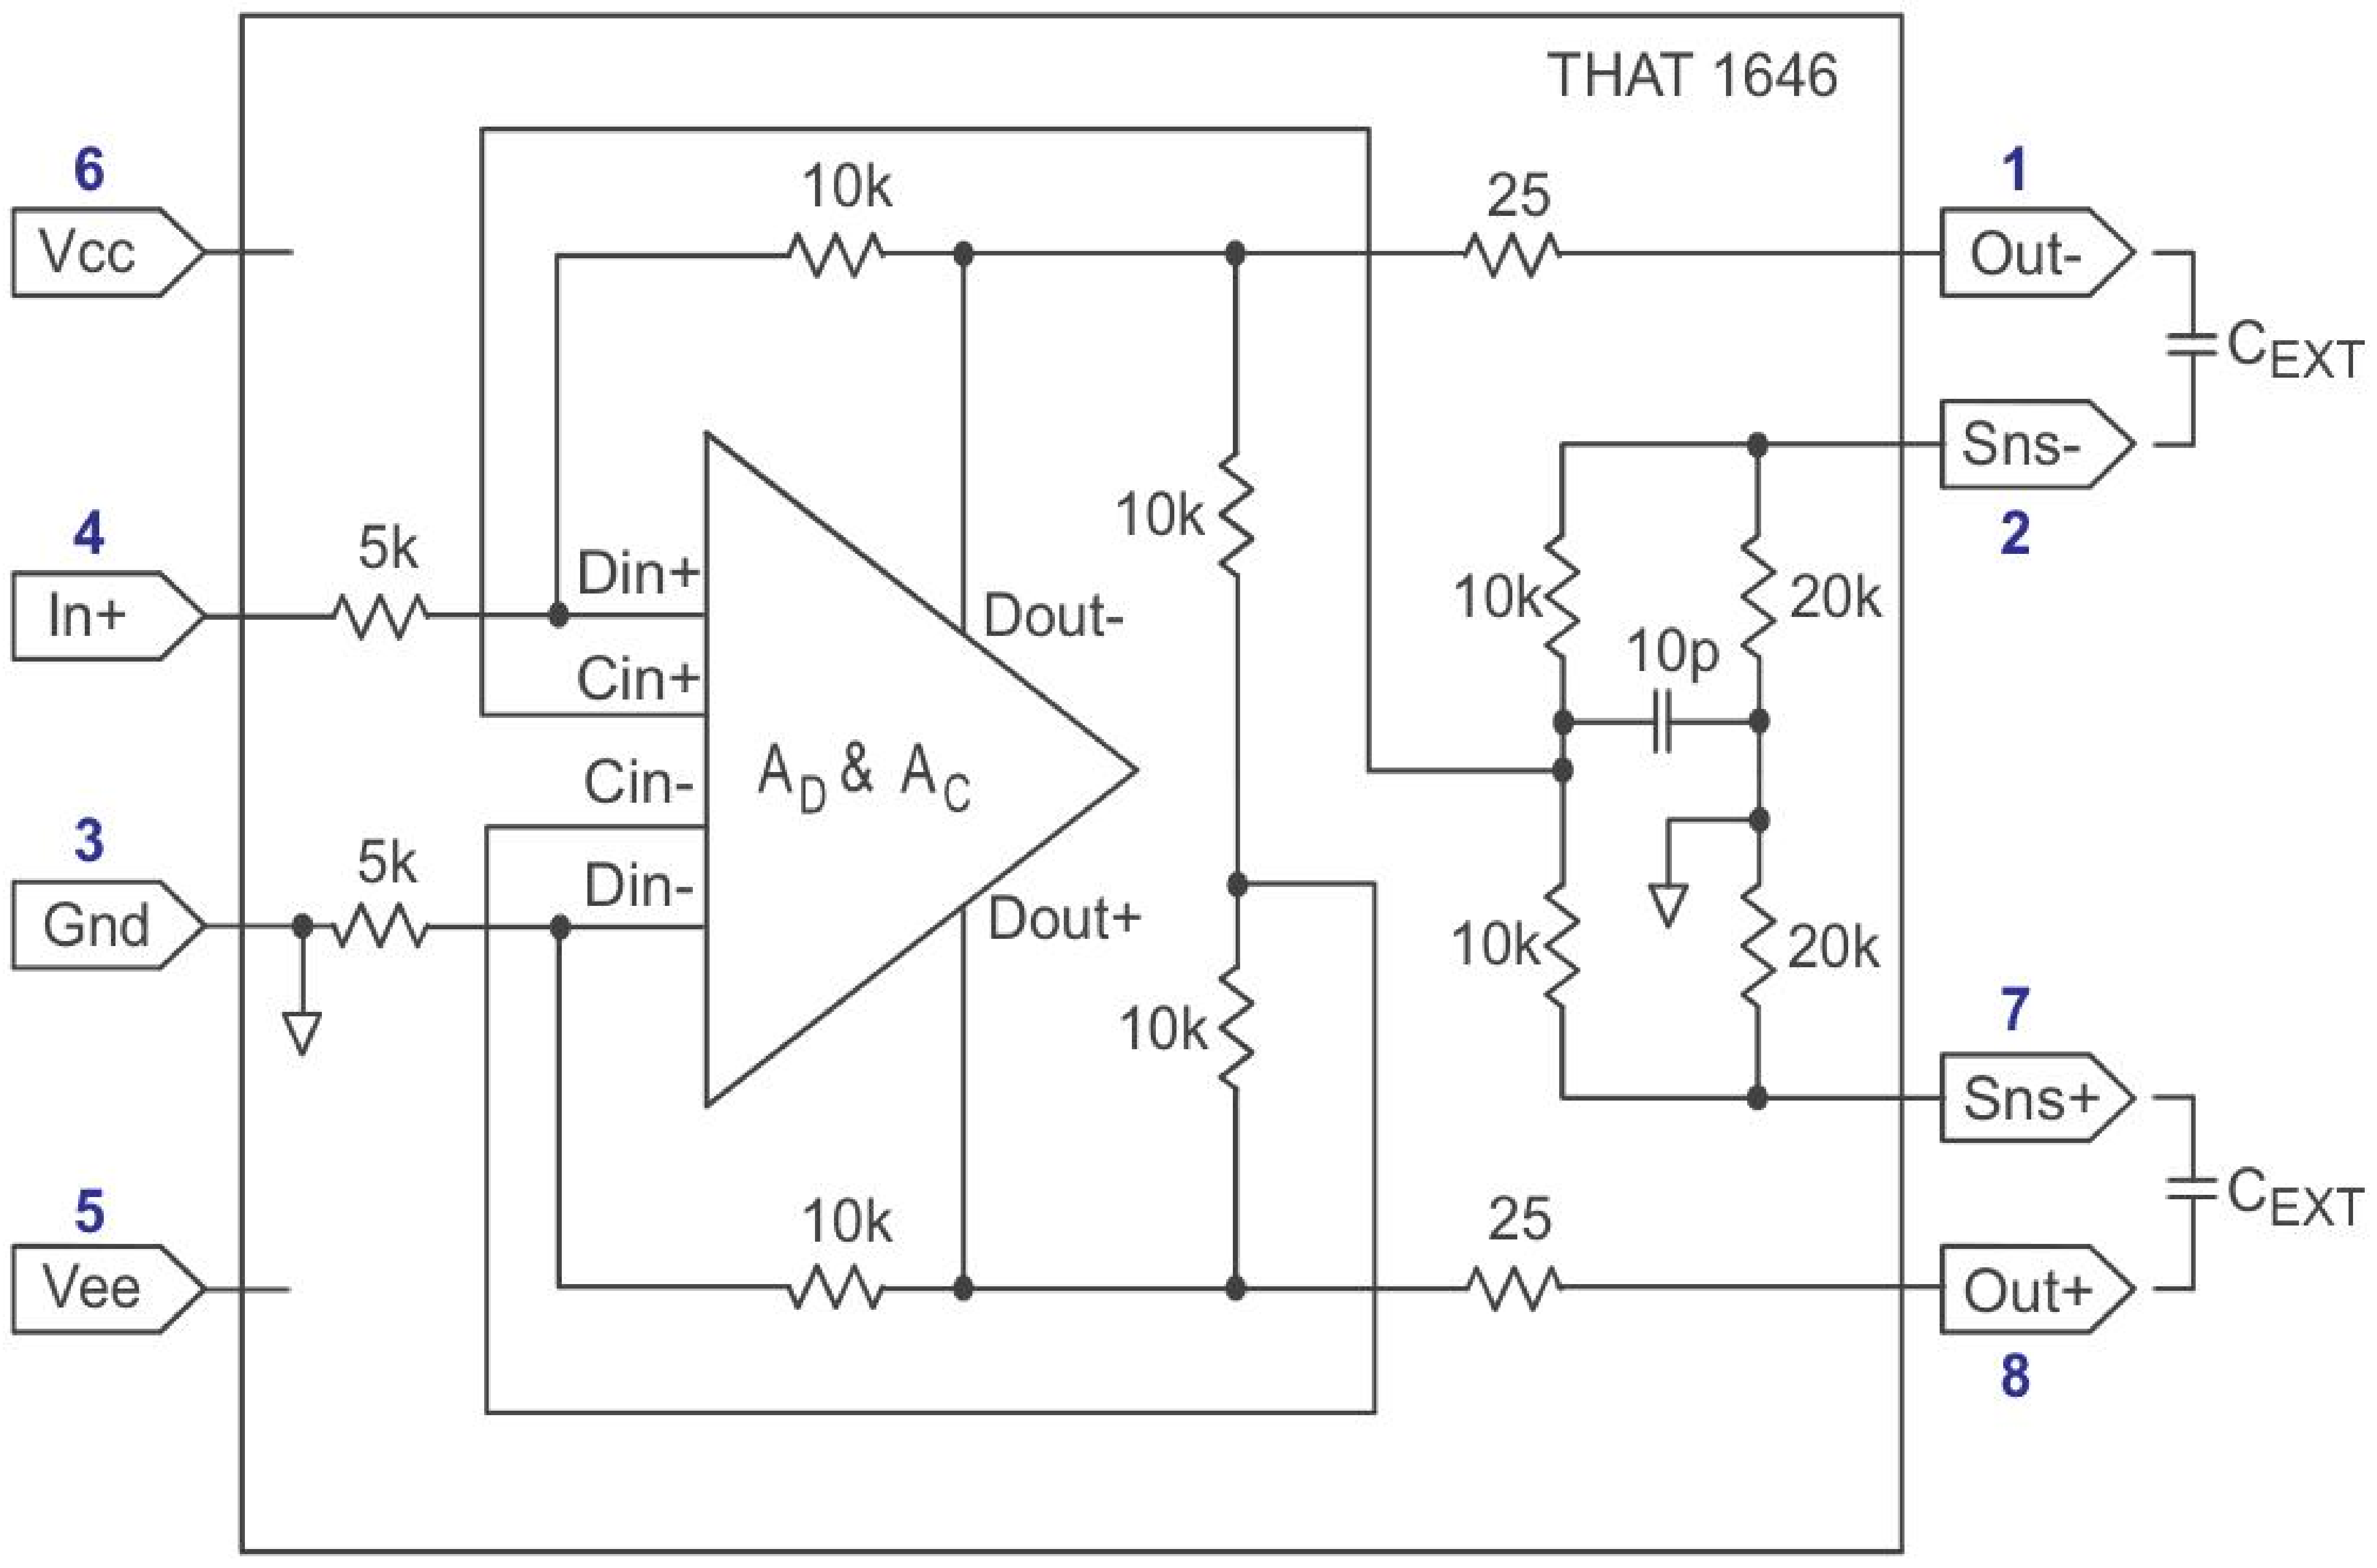

Fig.1: the circuit structure and pinout of the THAT1646 (from the device data sheet).

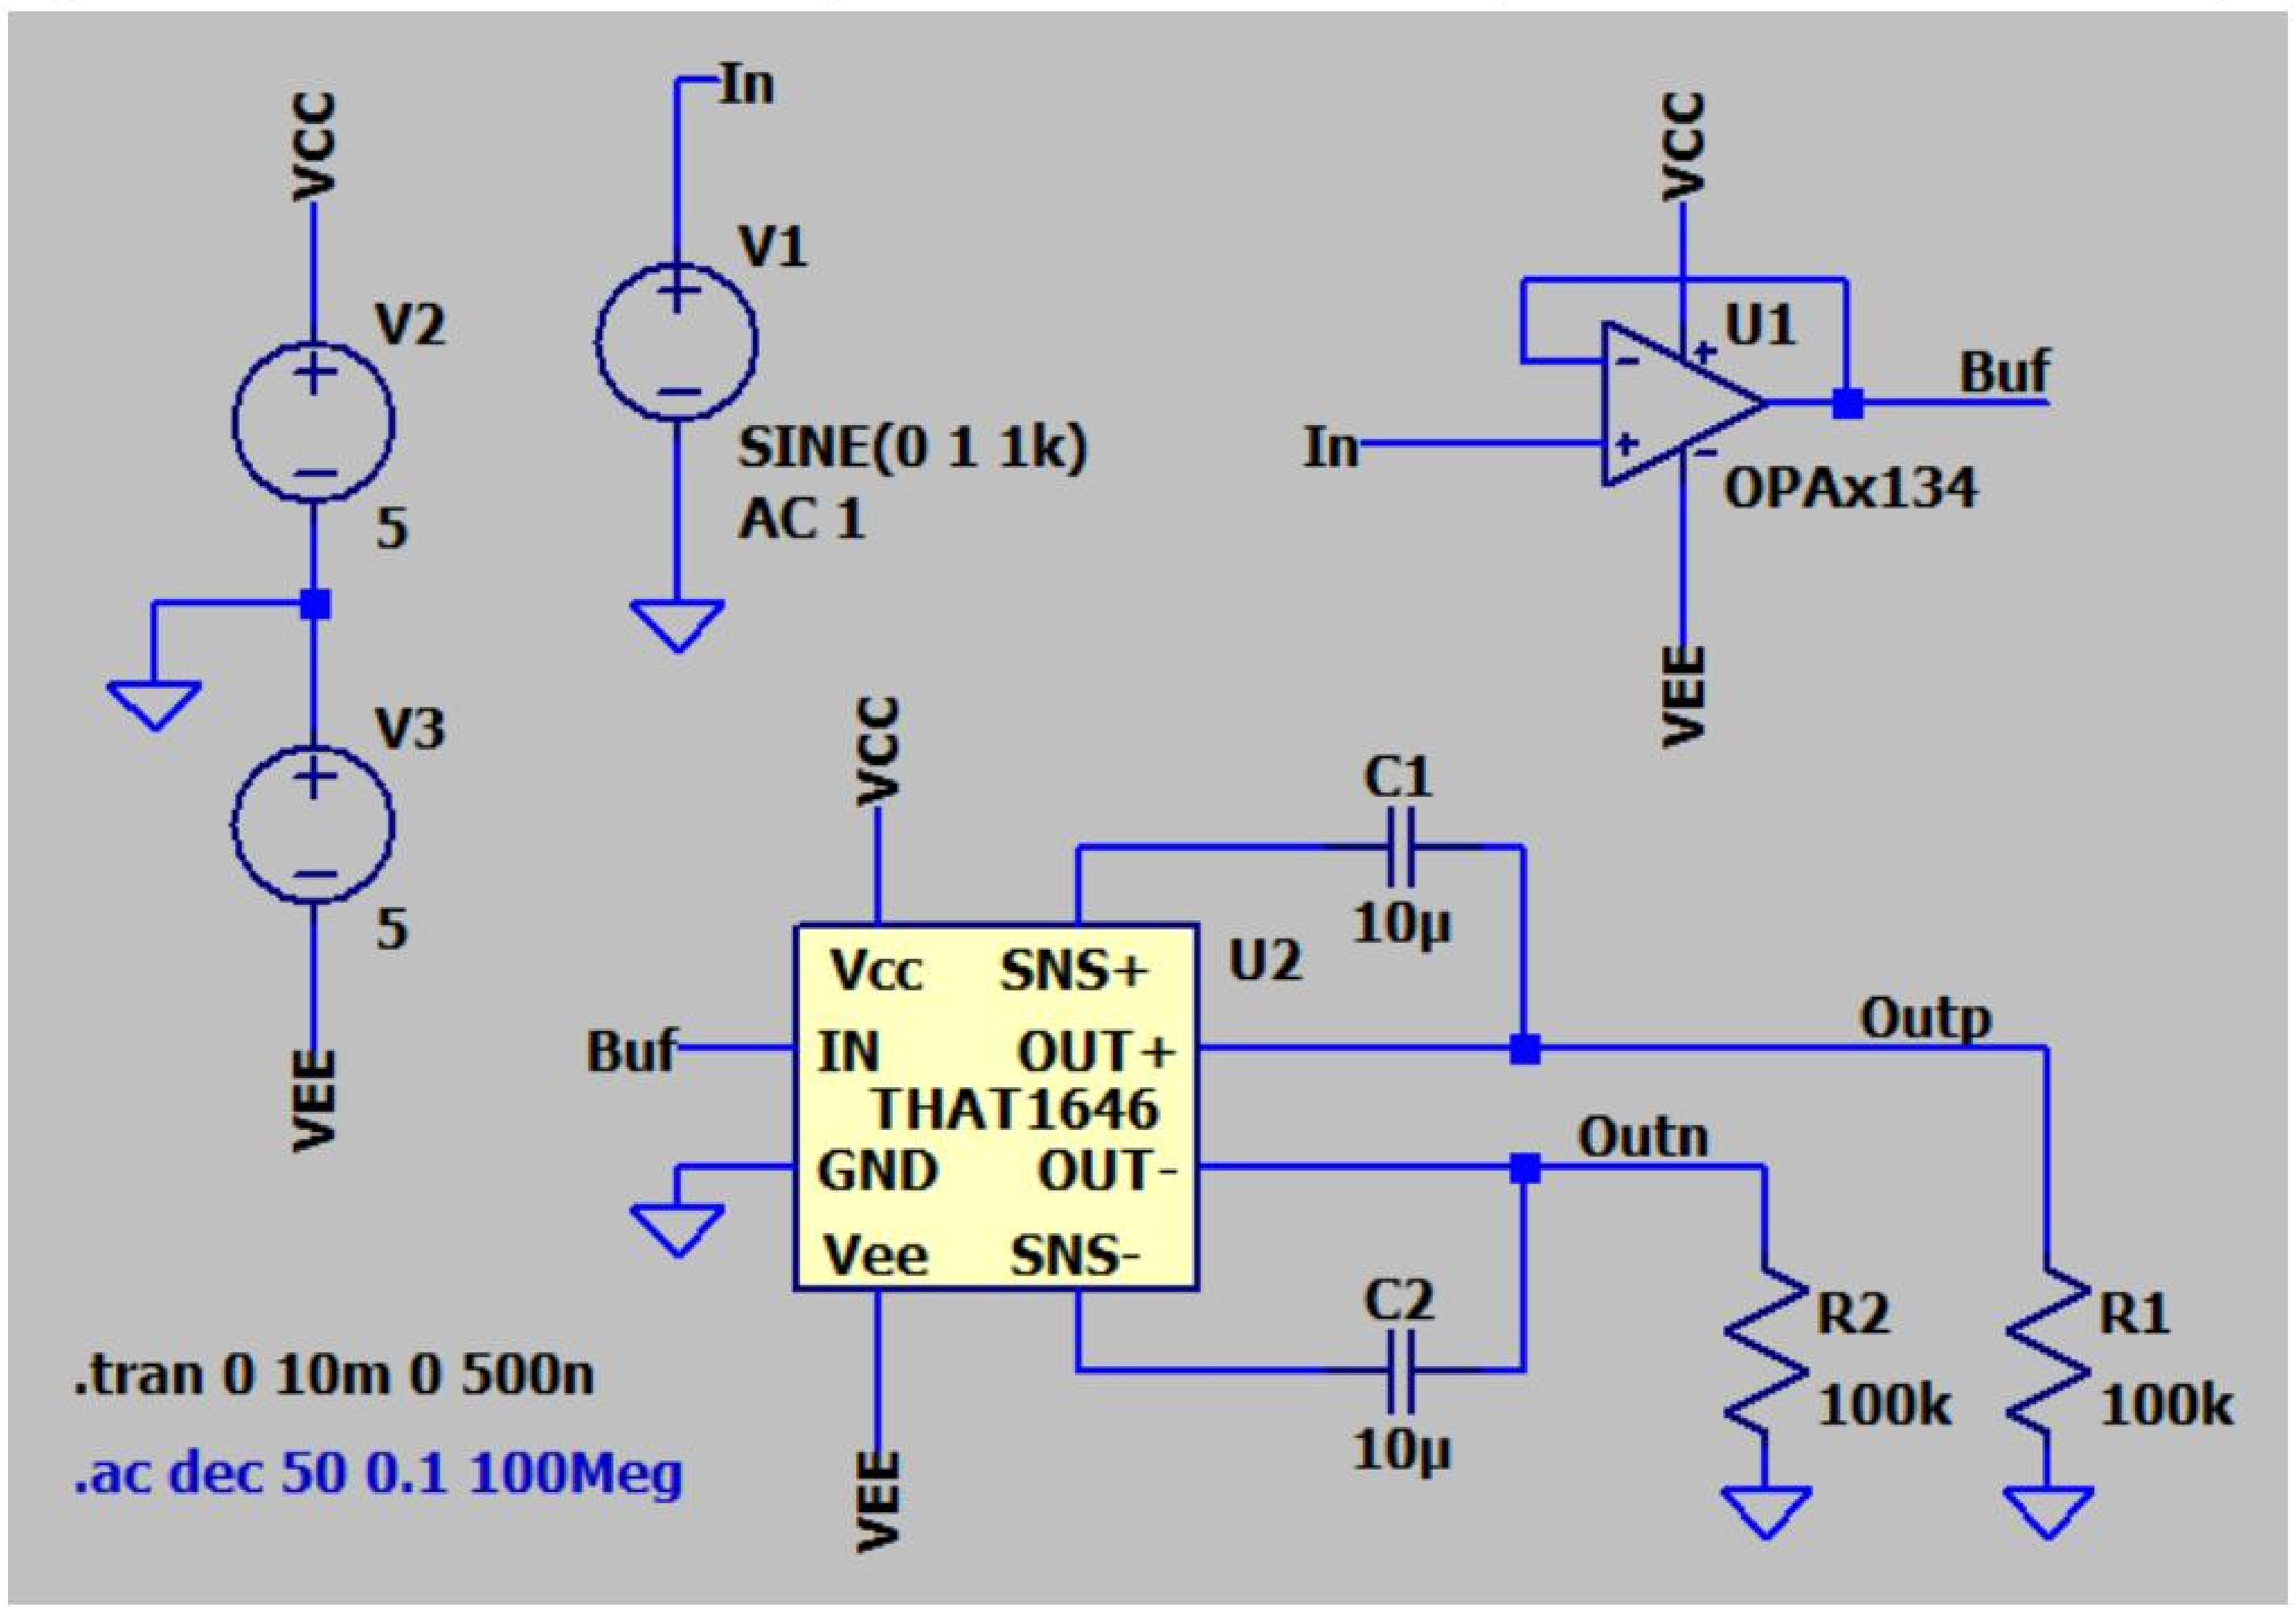

```
.tran 0 10m 0 500n
.ac dec 50 0.1 100Meg
```

Fig.2: an LTspice schematic for simulating the THAT1646.

as the THAT1646 has a relatively low input impedance of about 5kΩ.

Last month, we presented some transient simulation results but discussed concerns with a slow simulation rather than the results themselves. Fig.3 shows transient simulation results.

The top plane shows the single-ended input of the THAT1646 (the **Buf** signal from the buffer), the individual outputs are shown in the middle pane (green **outp** and yellow **outn** traces) and the differential output in the lower plane. The differential output trace is obtained by plotting **V(outp)–V(outn)** using the waveform maths.

All the panes in Fig.3 have the same voltage scale, so it is straightforward to see that the **outp** signal and **buf** signals are equal (1V peak sinewave), that is, the THAT1646's Out+ output has a unity non-inverting gain. Similarly, the THAT1646's Out- output (on the net **outn**) has unity inverting gain.

The differential signal thus has twice the amplitude of the two individual signals (this is the nature of differential signals). So, the gain of the THAT1646 to its differential output is 2. There is minimal phase shift (not visible on the timescale used in Fig.3).

Results from the AC analysis are shown in Fig.4. Unlike the transient simulation, this was obtained by driving the THAT1646 directly from the V1 source (signal **In** connected to the THAT1646's In pin). This is easily achieved by changing the net label on the wire connected to the THAT1646's input.

It is useful to draw LTspice schematics with subcircuits connected by name, rather than drawn wires, as this makes 'rewiring' to run alternative simulations easier. In LTspice, nets with the same name are connected; you do not always have to draw wires to connect components.

The THAT1646 was driven directly from **V1** for this simulation because an AC analysis is more idealised than transient analysis, and we are more interested in the device's raw characteristics than its performance in a typical circuit situation.

The bandwidth of the OPAx134 is fairly close to that of the THAT1646, so the OPAx134 will influence the overall frequency response if it is included.

You can easily run the analysis with and without the OPAx134 to compare them. You will find that the overall cutoff is sharper with the OPAx134 in place because it has slight gain peaking that coincides with frequencies just below the THAT1646's cutoff.

In Fig.4, the top trace shows the individual outputs and the lower trace the differential output, with the same colours as in Fig.4. The single-ended gains are very close to unity (0dB), while the differential gain is close to 2× (+6dB), at 5.97dB.

Note that the gain range plotted is small (-1dB to +1dB), which allows us to see a small variation in gain of the individual outputs (a difference of about 0.2dB). This variation is symmetrical, so it is cancelled and does not appear in the differential plot.

As mentioned last month, differential circuits can cancel some of their own imperfections as well as being an essential component in interference-cancelling balanced connections.

From Fig.4, we can observe that the gain and phase responses are both very flat to well over 1MHz. A look at an expanded gain range (to see the -3dB points) indicates a bandwidth of around 10MHz, which is consistent with the datasheet value.

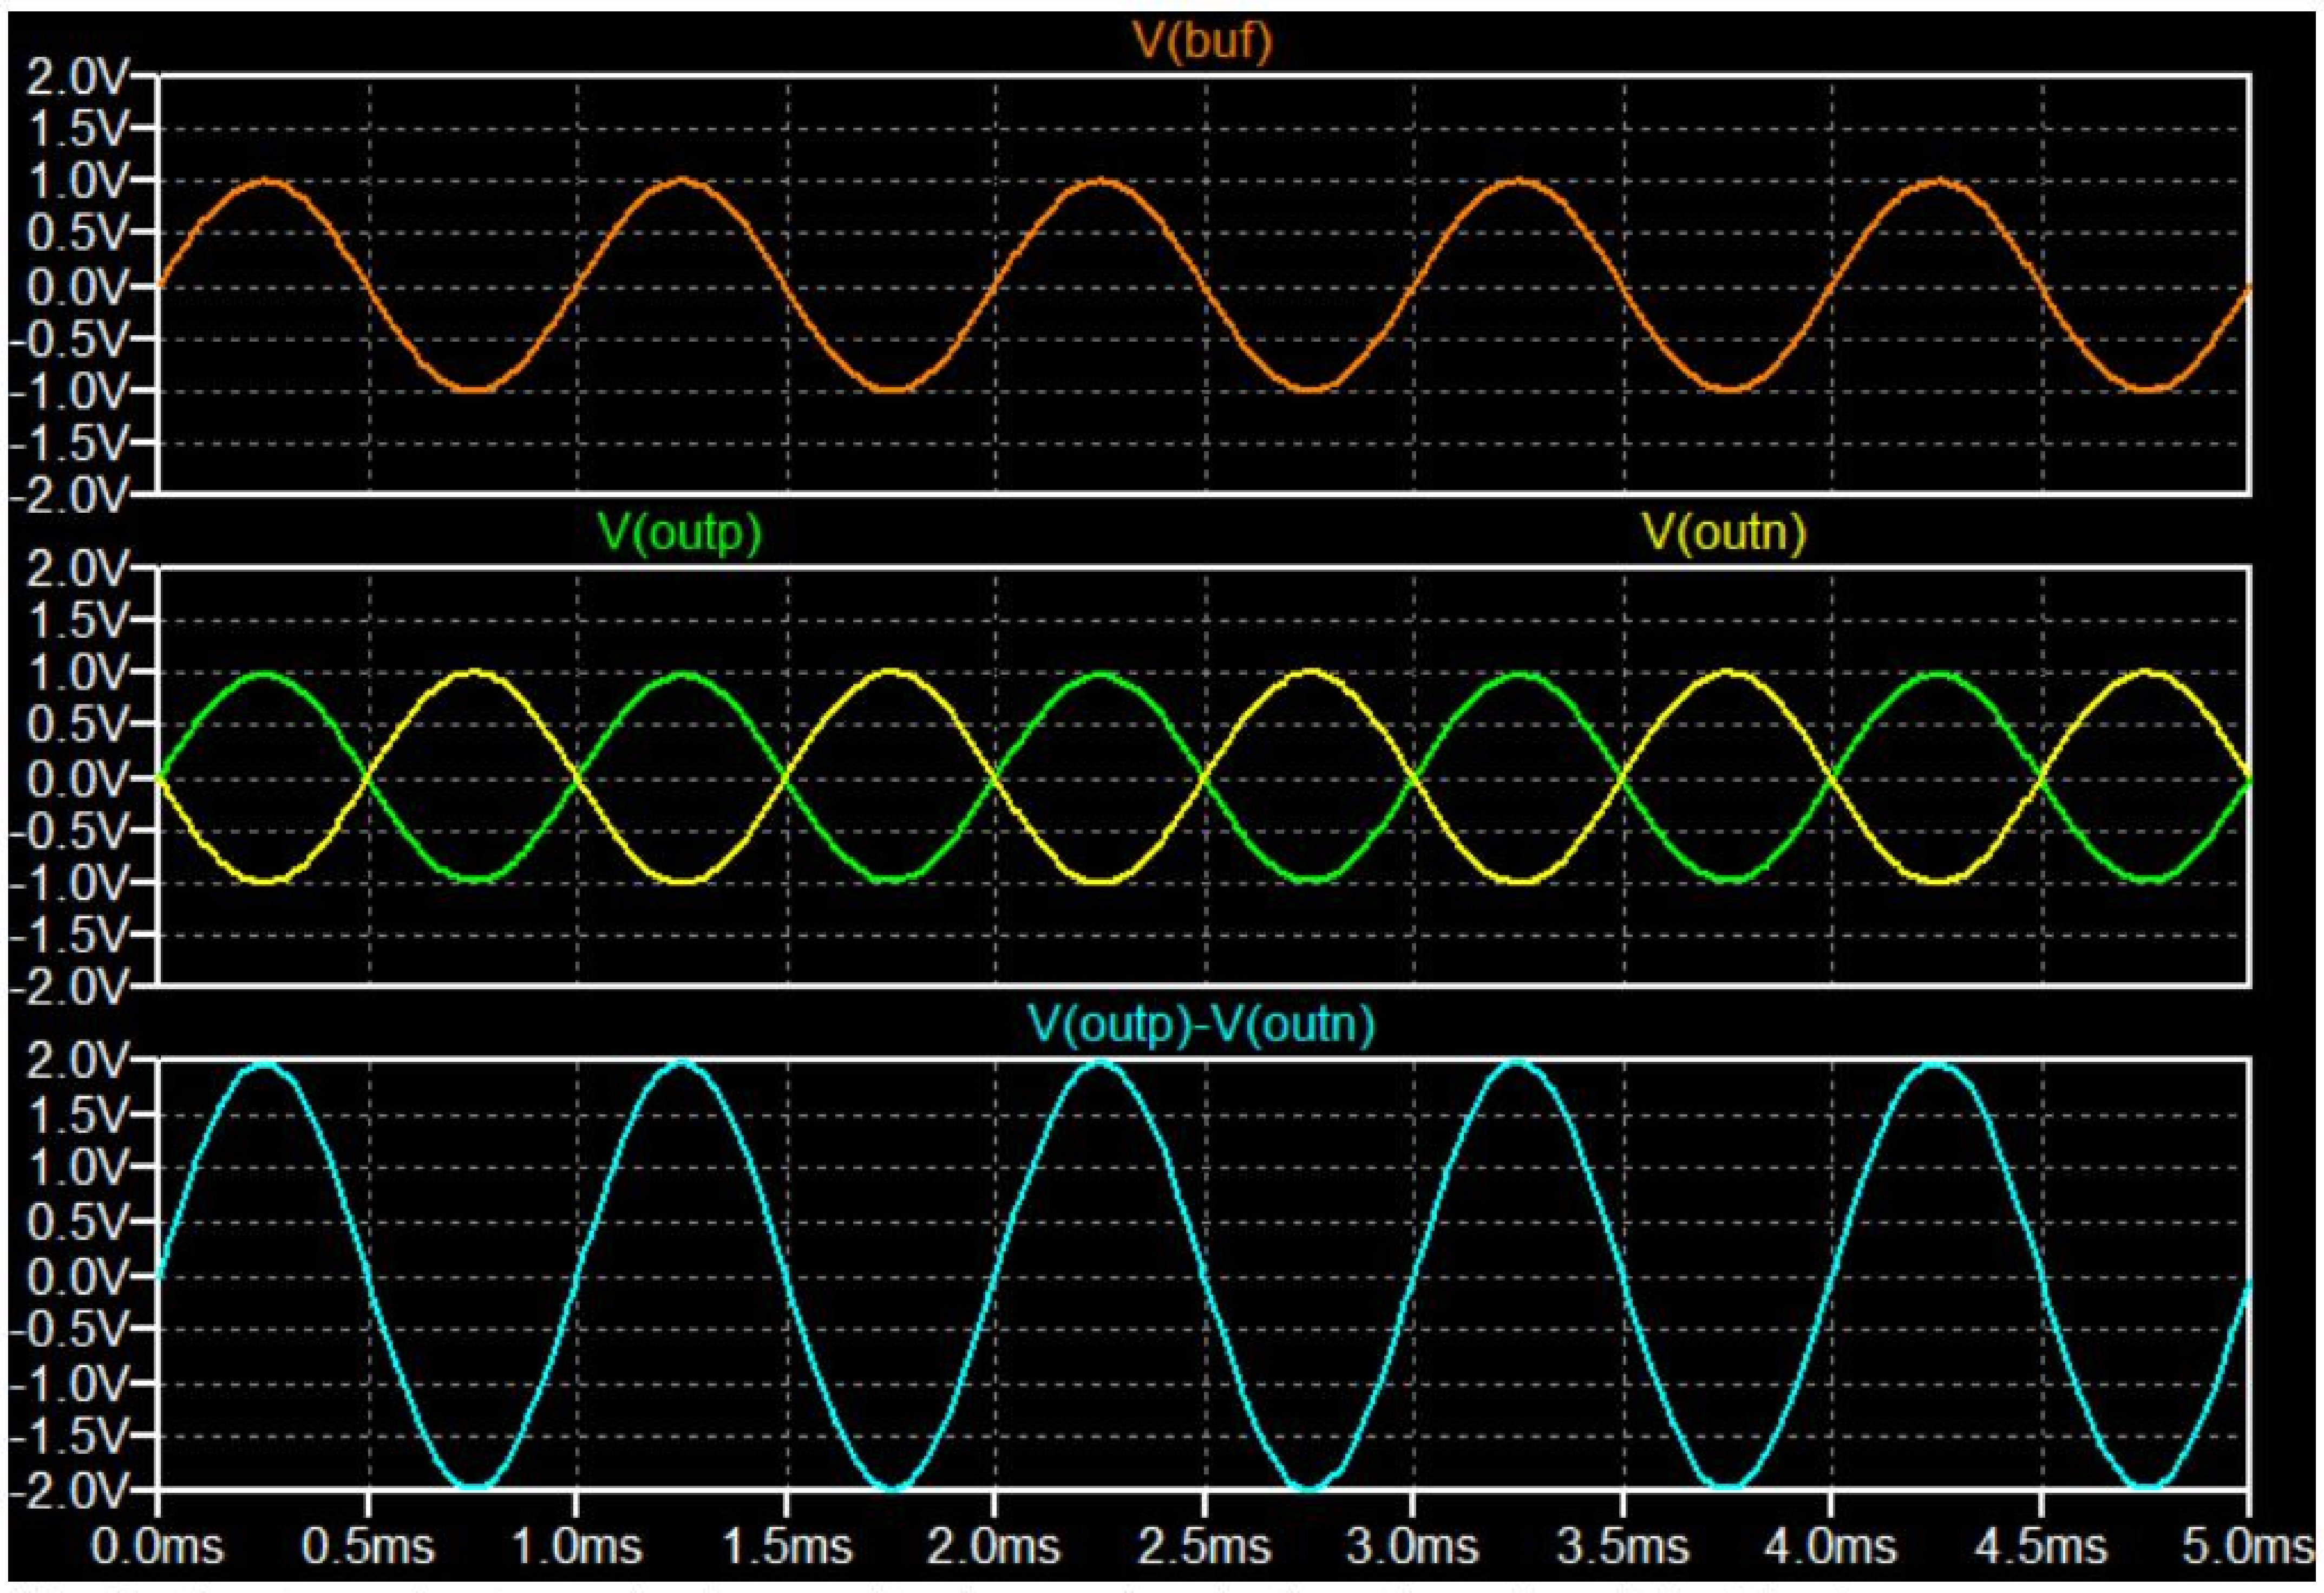

Fig.3: the transient analysis results from simulating the circuit in Fig.2.

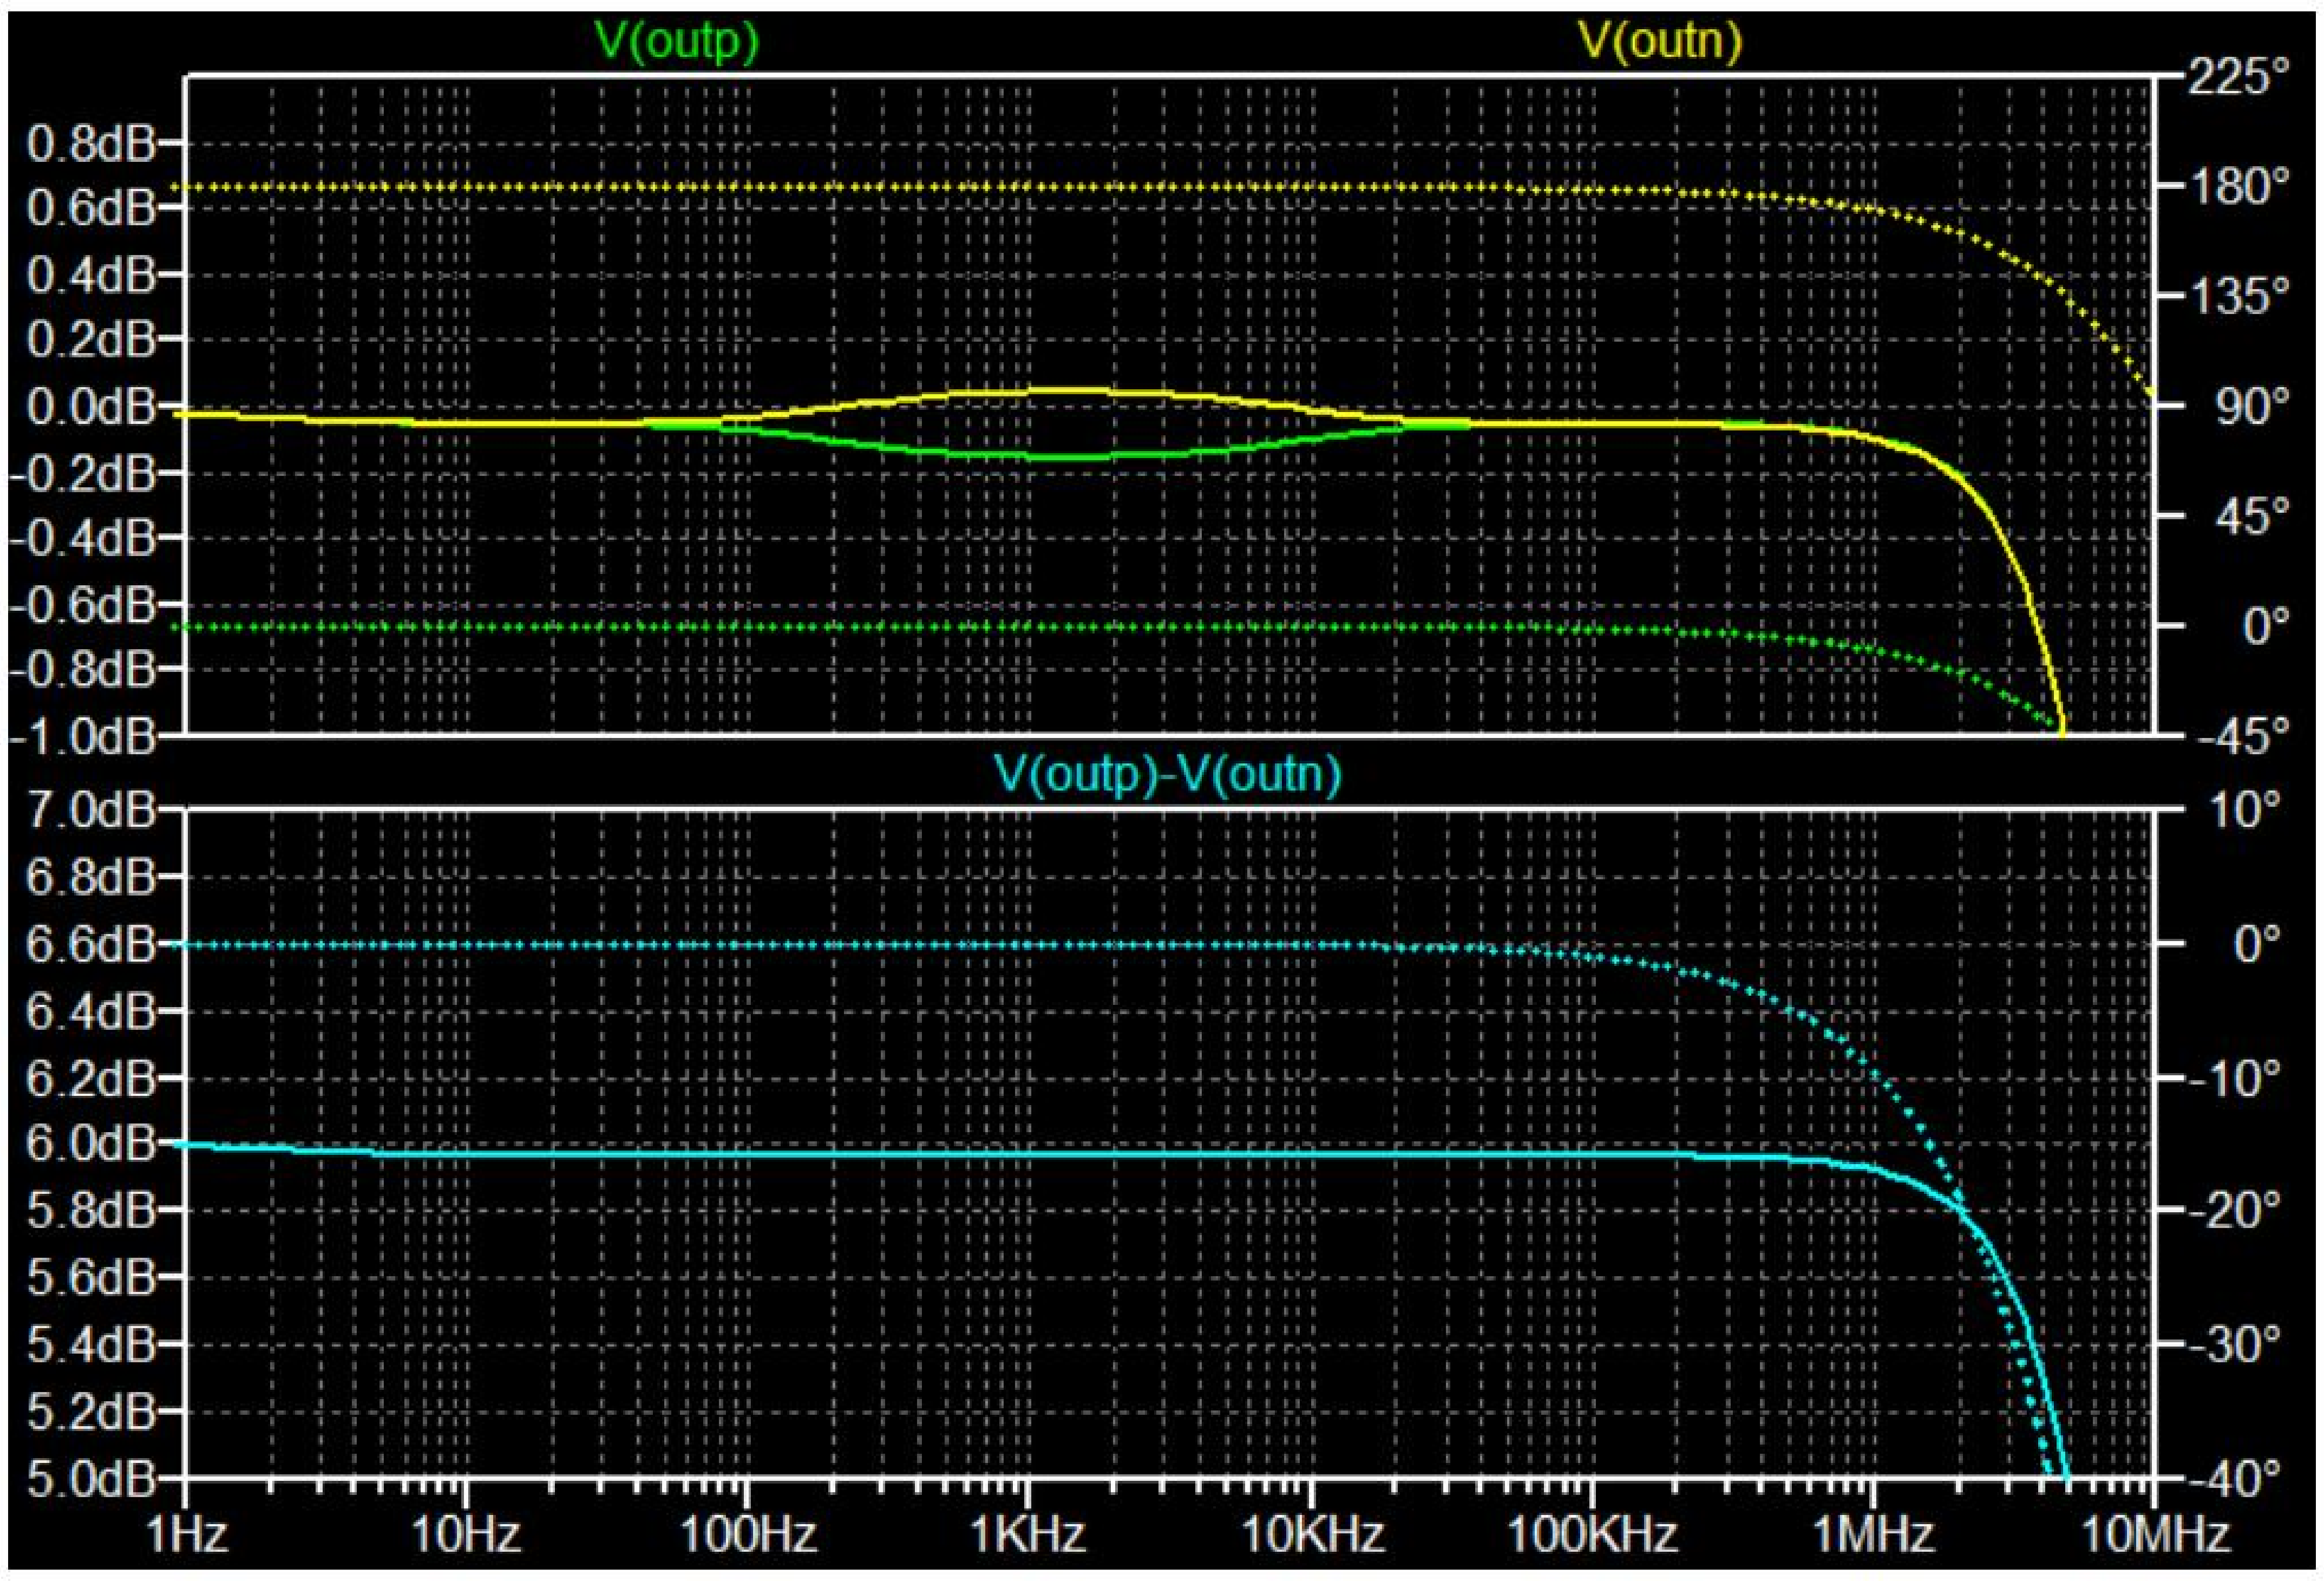

Fig.4: the AC analysis results from simulating the circuit in Fig.2 with R1=R2=100kΩ.

5-year collections

2019-2023

£49.95

2018-2022

£49.95

2017-2021

£49.95

2016-2020

£44.95

Purchase and download at:
www.electronpublishing.com

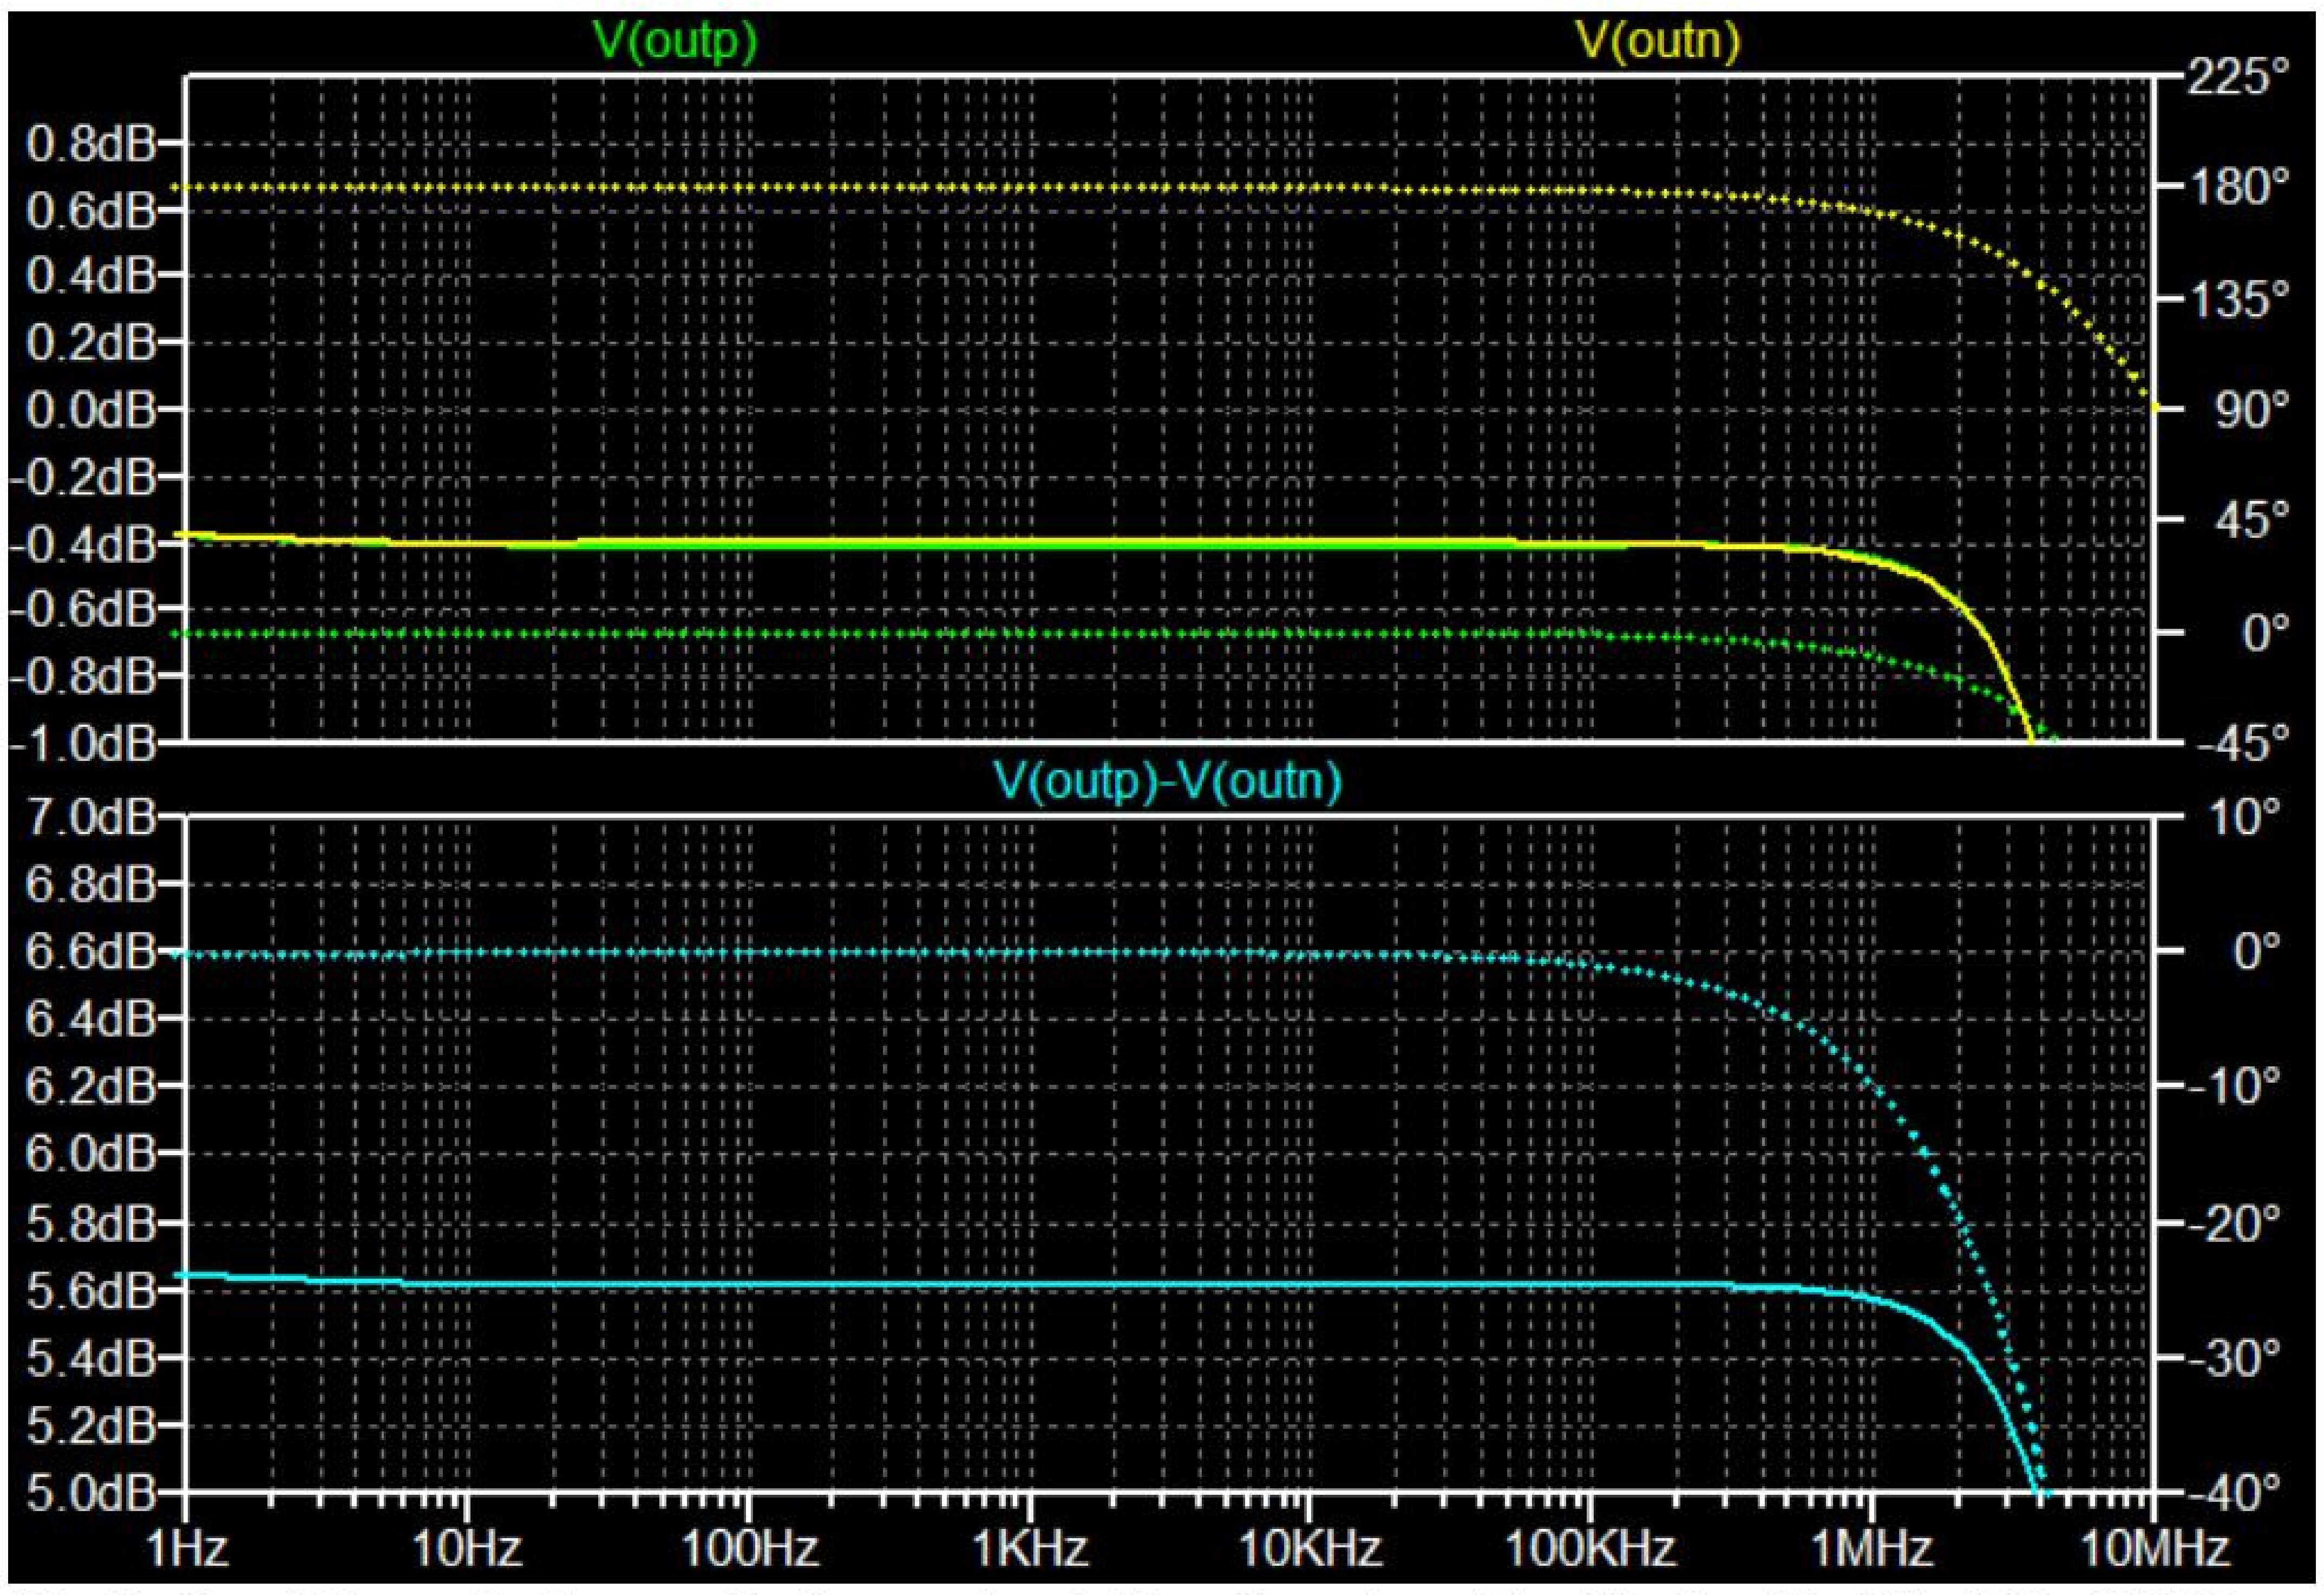

Fig.5: the AC analysis results from simulating the circuit in Fig.2 with R1=R2=600Ω.

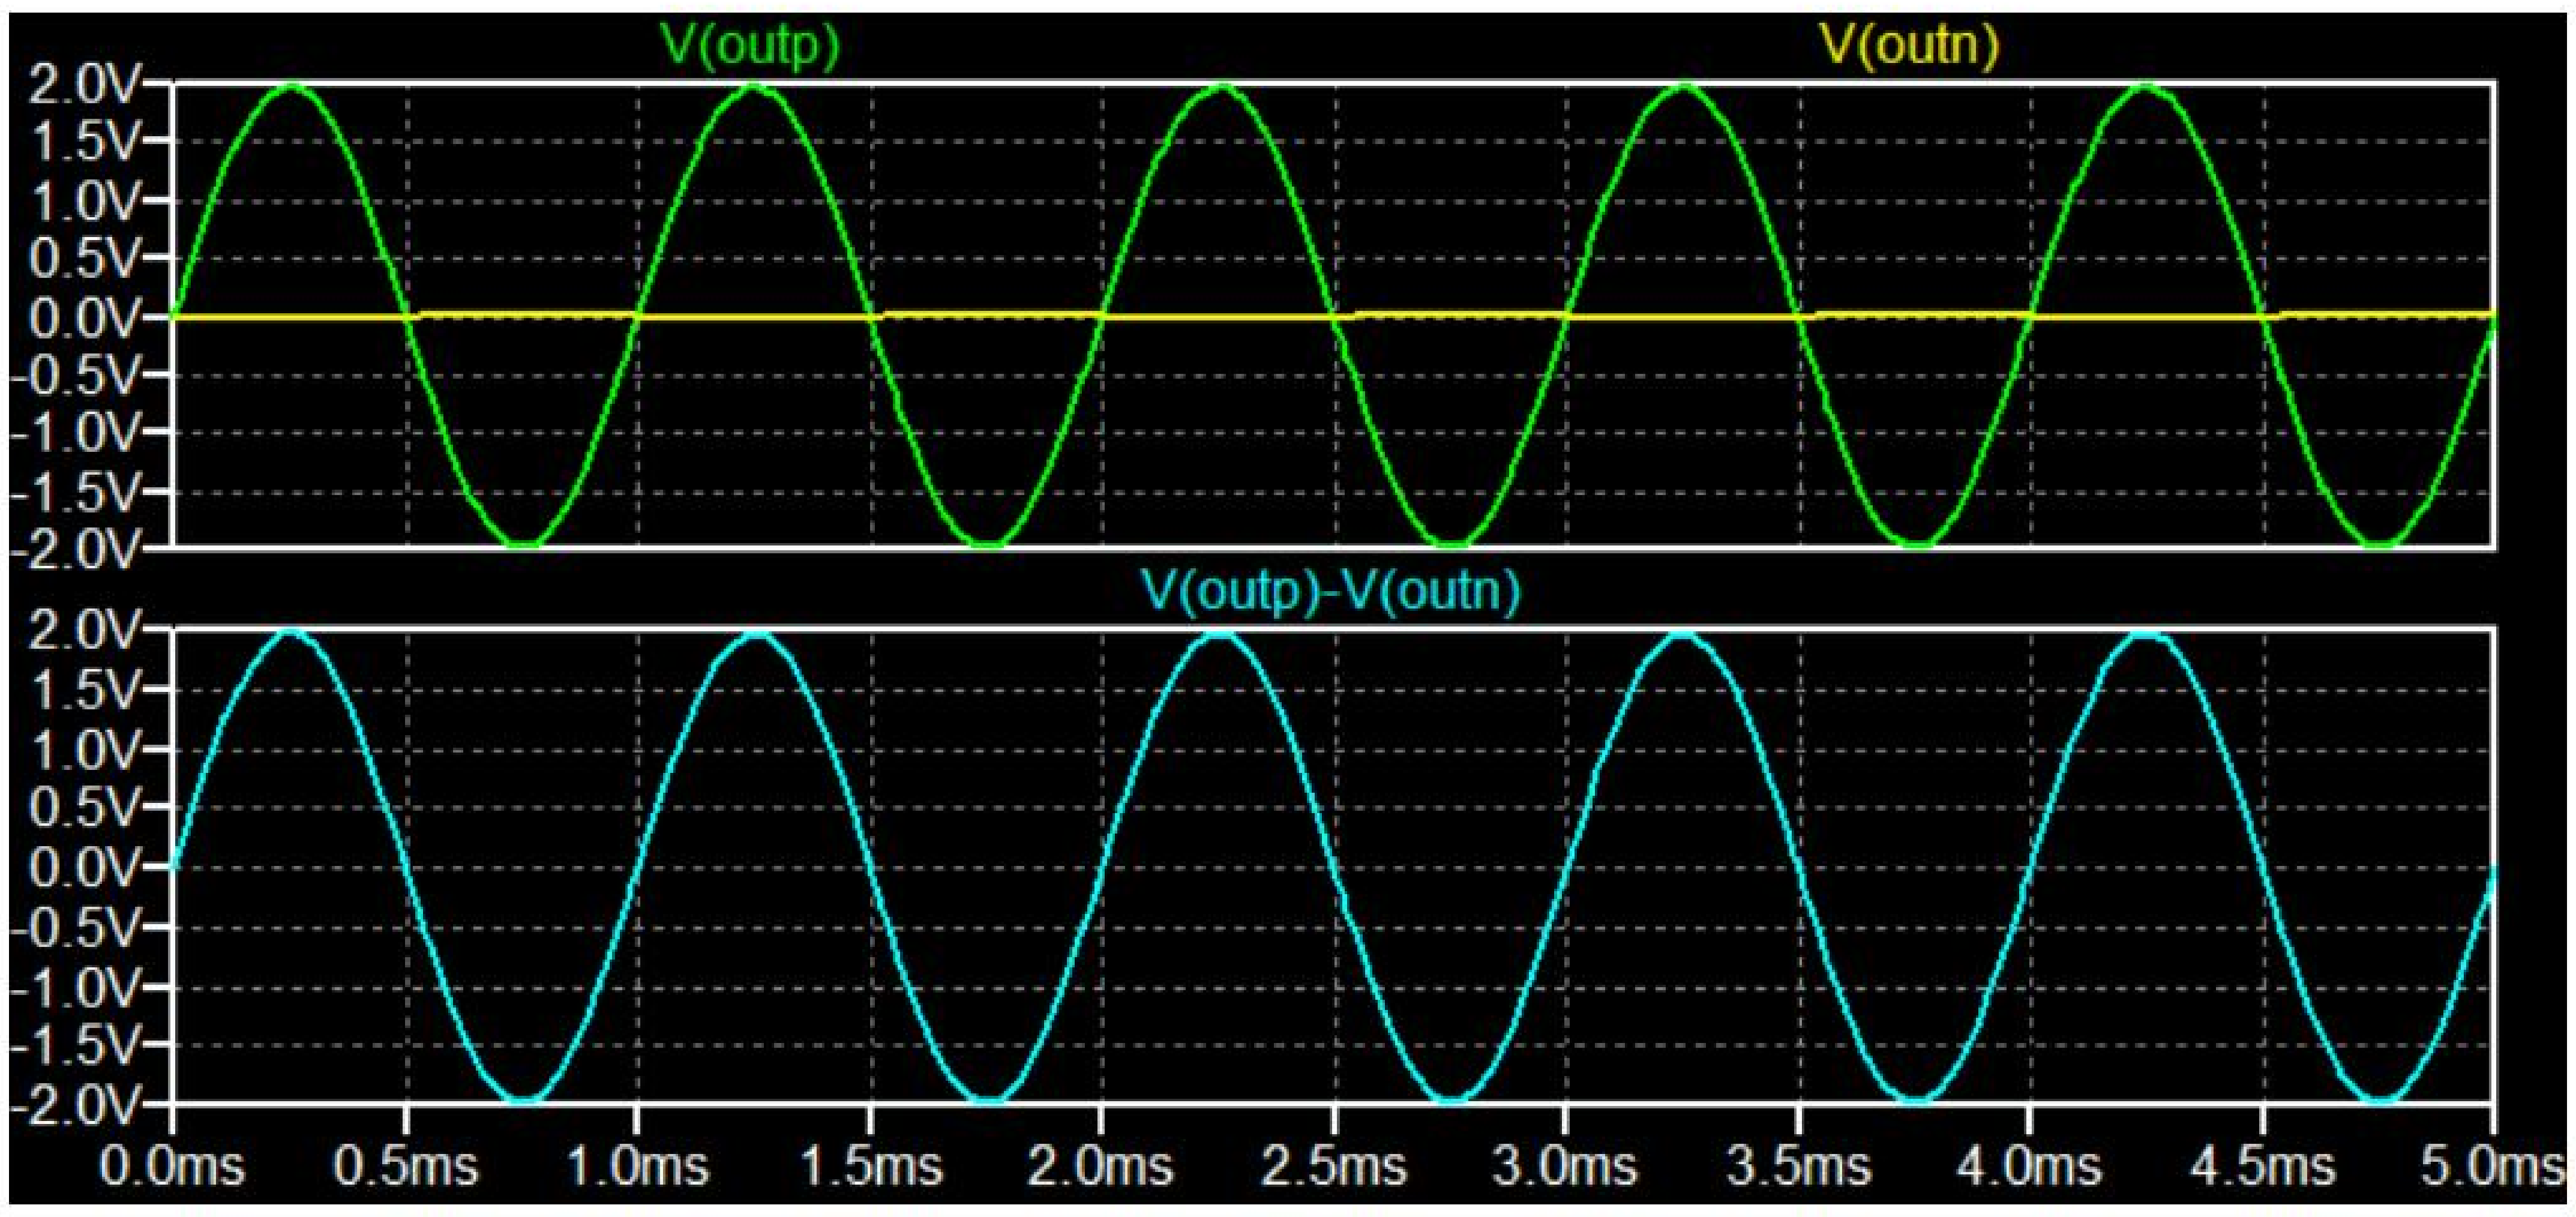

Fig.6: the transient analysis results of the circuit in Fig.2 with **outn** shorted to ground.

Fig.5 shows the AC simulation with the load resistors reduced from 100kΩ to 600Ω. The gain has dropped by 0.35dB from 5.97dB to 5.62dB. This is due to the higher loading presented by the 600Ω resistors.

Fig.6 shows transient simulation results with the one output (**outn**) shorted to ground. This was achieved by setting the value of **R2** to 1mΩ (**R1** is 100kΩ, as for the result in Fig.3).

It is convenient to do this rather than connecting the THAT1646's OUT- pin directly to ground as it involves fewer circuit changes, is more realistic and means we can keep the same plot setup (if we short the output, the net **outn** no longer exists).

The input is the same as previously (a 1V peak 1kHz sinewave), and despite one output being shorted, the differential output (lower pane in Fig.6) is the same as in Fig.3 (2V peak), showing that the gain is unchanged.

The upper pane shows that **outn** is at 0V as expected, and **outp** is 2V peak, rather than the 1V without the short circuit. This shows that the THAT1646 is behaving as a floating output, the same behaviour as exhibited by an output transformer, as discussed last month.

Input and output impedance

The loading of one circuit (the source) by another (the load) is a common challenge in electronics. The source may be the output of a circuit, such as an amplifier or filter, or it may be a sensor, such as a microphone or photodiode. It could also be a power source, like a battery.

The load may be the input of a circuit such as an amplifier or filter, a transducer such as a loudspeaker or motor, or the circuit or device connected to a power supply.

The output of any circuit between two terminals can be represented by an ideal output voltage source (v_O in Fig.7) in series with an output (or internal) impedance (Z_O in Fig.7). The load can be represented as a single impedance (Z_I Fig.7). The voltage (signal) between the load and the source is shown as v_s in Fig.7.

The term output impedance is usually used for signal sources, while internal impedance is used for power sources. In many cases, for both the source and input, we can use resistance instead of impedance to simplify the situation, particularly if there is little frequency dependence over the range of interest.

The output impedance of a circuit may be dominated by obvious components such as an output resistor or coupling capacitor, but often it is a complex function of the properties of the circuit as a whole. The representation in Fig.7 is an equivalent circuit, which is a simple model of how the circuit behaves.

The source equivalent circuit is based on Thévenin's theorem and strictly only applies to linear circuits, but we often simplify situations by assuming linearity. For example, AC analysis is based on linearised circuit models. We can also model an output as a current source in parallel with an impedance or resistance (a Norton equivalent circuit) if that is more convenient.

In Fig.7, we see that the input and output impedances act as a potential divider, reducing the signal voltage on the connection (v_s) with respect to the internal source output (v_O). Using the potential divider formula $v_s = Z_I v_O \div (Z_I + Z_O)$. This is the loading effect – for maximum signal amplitude transfer, we need Z_I to be large and Z_O to be small.

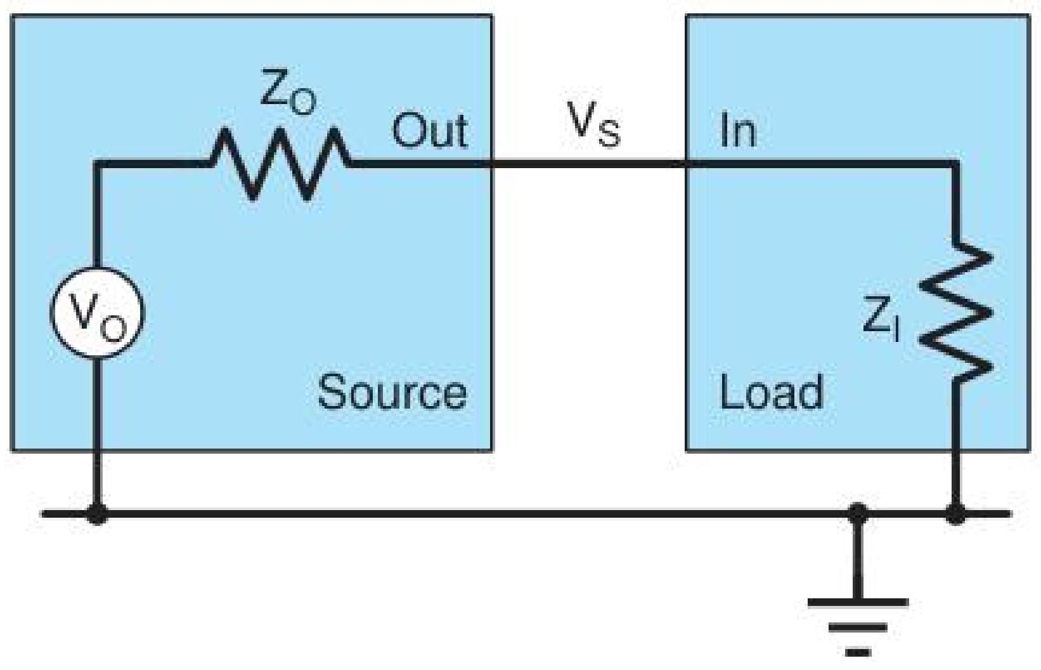

Fig.7: a generic source and load equivalent circuit.

If we have $Z_I = Z_O$, the connection is matched. The voltage is attenuated by half, but we get maximum power transfer. Matching is required where the connection is more than about 1/10th the wavelength of the signal, but this is in the order of kilometres at audio frequencies, so it is rarely a concern.

There is some legacy of 600Ω matching from old telephony systems in audio, but usually we need maximum amplitude transfer for signal connections.

An equivalent circuit for a single-ended connection with a driver output impedance (Z_O) and receiver input impedance (Z_I) is shown in Fig.8. This is very similar to Fig.7, but we have added the receiver amplifier symbol as a reminder of how the input is connected. There is also an interfering voltage source in the

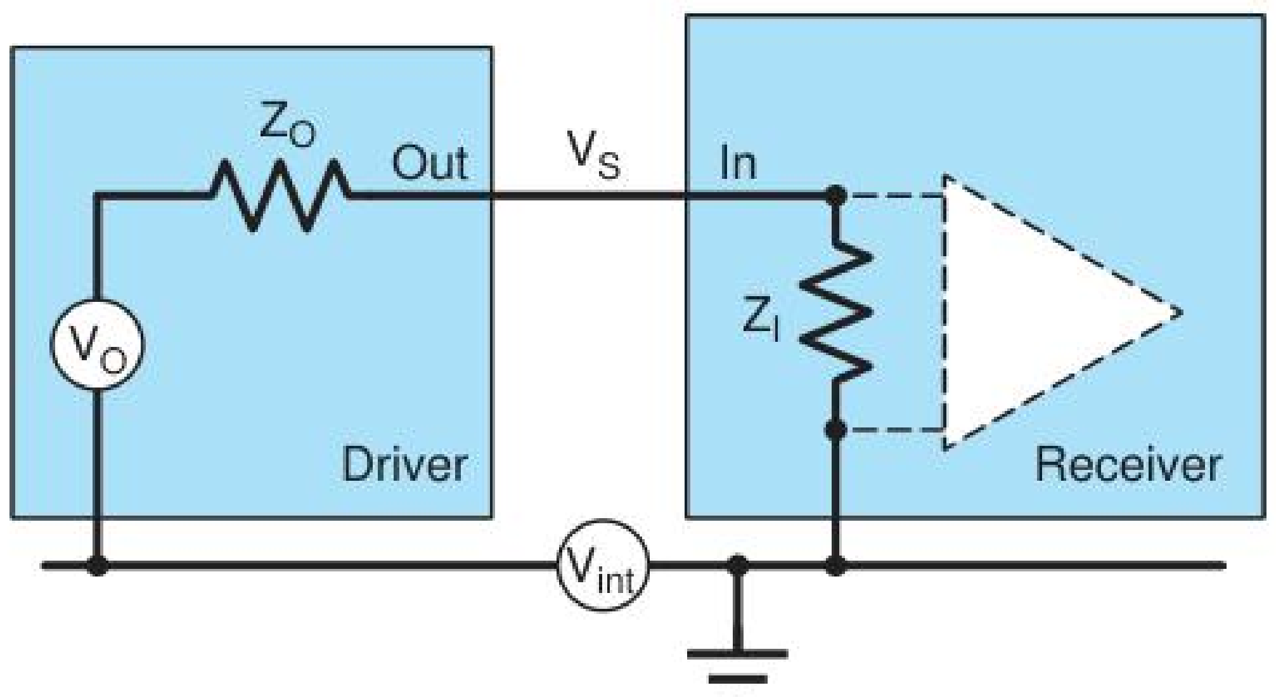

Fig.8: source and load equivalent circuits for a single-ended line driver & receiver.

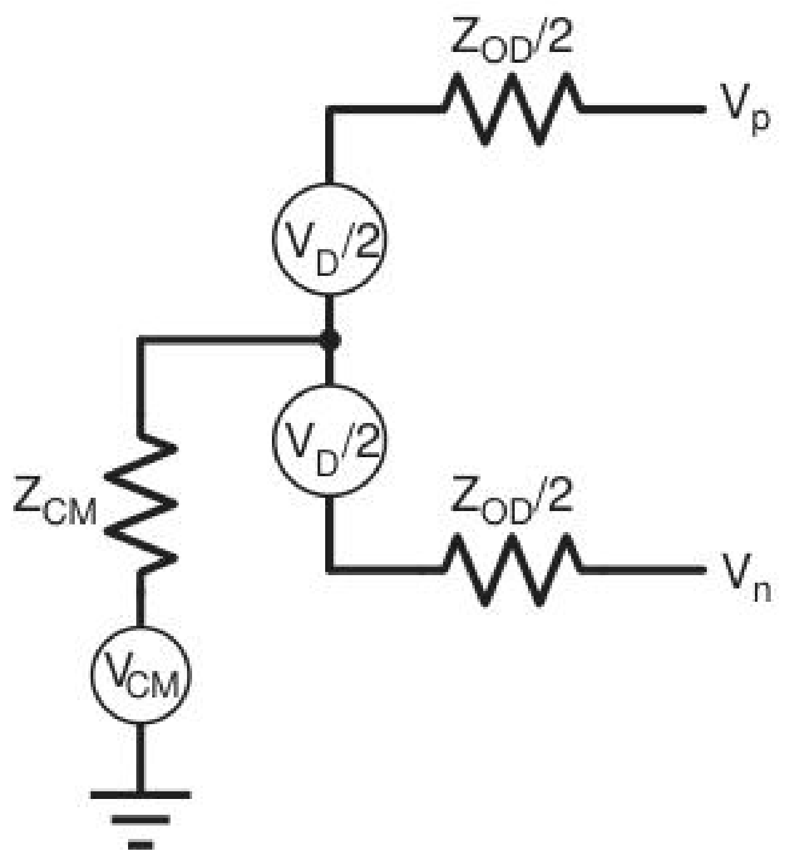

Fig.9: a possible differential plus common-mode source model.

ground connection between the driver and receiver (v_{int}).

We discussed interfering signals in depth last month – balanced lines can provide good interference rejection, whereas this is not possible for single-ended lines unless the frequency ranges of the wanted and unwanted signals have no overlap. In Fig.8, we see that the interfering signal is in series with the source output, so effectively they behave as one source.

Note that we use the receiver ground as the reference point, as we are interested in the integrity of the signal at the receiver.

Balanced line equivalent circuit

The equivalent circuit for a balanced line driven by a differential output is more complex than Fig.8. Last month, we discussed common-mode and differential signals, and it follows that we can have both differential and common-mode input and output impedances.

Fig.9 shows a possible representation of these impedances based on the differential plus common-mode source model discussed last month. In Fig.9, the differential output impedance (Z_{OD}) is split to match the split in the differential source (v_D), but like the basic model in Fig.7, it is in series with the sources. Similarly, the common-mode source has a series impedance (Z_{CM}).

The common-mode source and impedance in Fig.9 naturally relate to the common-mode signal produced internally by a line driver amplifier. If Fig.9 represents a fully differential amplifier output, then v_{CM} would reflect its inability to perfectly reject common-mode inputs, so we could refer to this as input-related common-mode rejection.

Input common-mode rejection is important, but our main concern in our recent discussion of balanced lines was interfering signals that directly couple to v_p and v_n as common-mode signals external to the driver.

Fig.10 shows an equivalent circuit for a floating output differential/balanced driver and receiver connection, which includes common-mode signals introduced at the output.

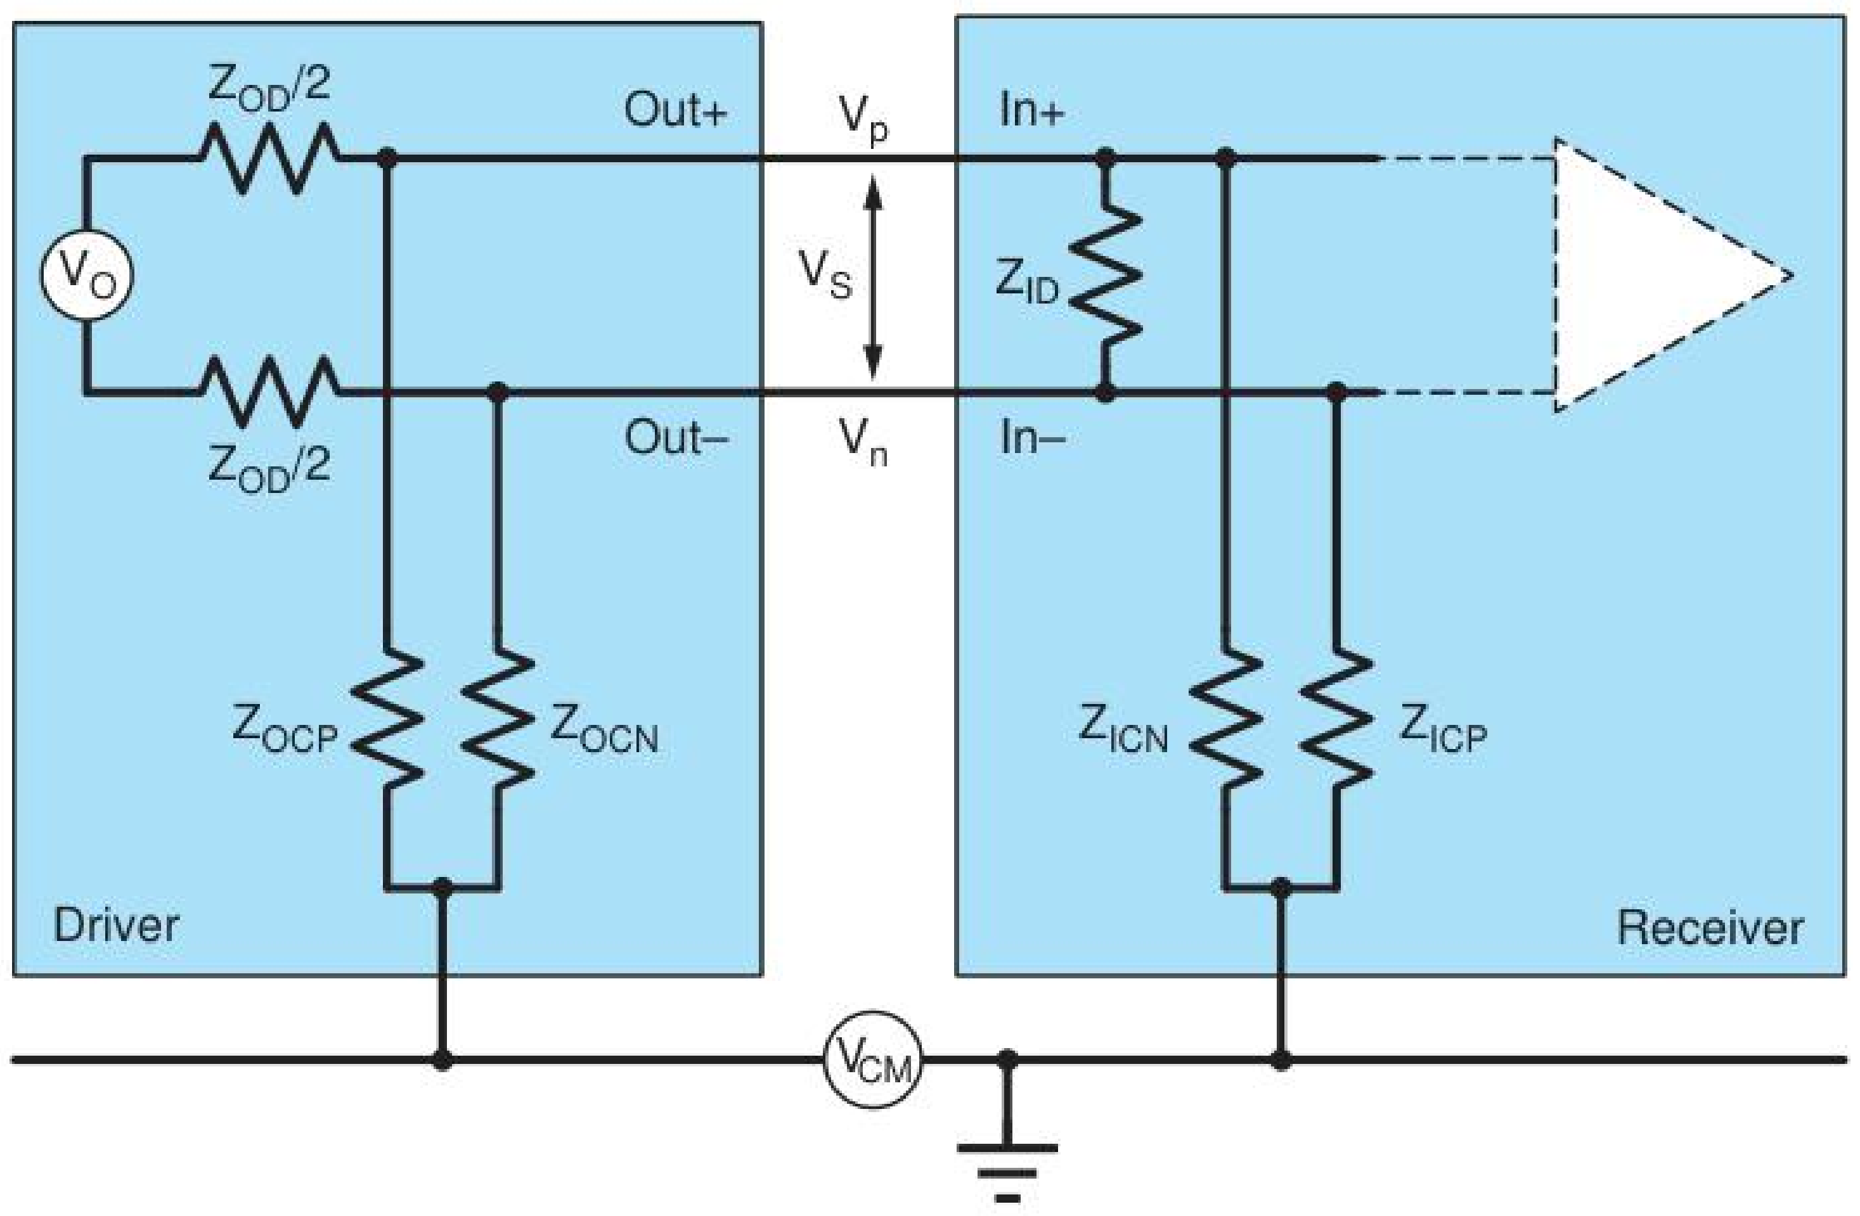

Fig.10: a source and load equivalent circuit for a balanced/differential line driver and receiver.

In Fig.10, the driver has common-mode output impedances to ground from each output, while the receiver has common-mode input impedances to ground from each input. These form two parallel potential dividers across the common-mode signal.

If the two output and two input common-mode impedances are both equal ($Z_{OCP} = Z_{OCN}$ and $Z_{ICP} = Z_{ICN}$), the line will be balanced and the voltage contributions to v_p and v_n from the common-mode source will be equal.

As discussed last month, this will minimise the interference picked up by the receiver, assuming the receiver is able to reject common-mode signals (more on common-mode rejection later).

From the perspective of the differential output, the driver common-mode output and receiver common-mode input impedances are in parallel. For the results in Fig.5, we can assume that the 600Ω load resistors are small enough to dominate the parallel value, so we can ignore Z_{OCP} and Z_{OCN}, and thus we have $Z_{ICP} = Z_{ICN} = Z_{CM} = 600\Omega$.

If Z_{ID} is large compared to the other impedances, it will have relatively little impact (note that in the Fig.2 circuit, Z_{ID} is not present).

With these assumptions, we can draw an equivalent circuit for the differential output as shown in Fig.11. This does not include any common-mode signals; we are just considering the differential component here.

We see that $Z_{OD}/2$ and Z_{CM} form potential dividers with respect to both halves of the differential output. The two halves of the differential signal will be attenuated by the potential divider factor, so the differential signal as a whole will be attenuated by the same factor.

From the THAT1646 datasheet, its output impedance is 40-60Ω. The two 25Ω resistors in Fig.1 correspond to the $Z_{OD}/2$ resistors in Fig.9 for a nominal output impedance of 50Ω. The attenuation seen in Fig.5 is 0.35dB, which is a factor of $10^{(0.35 \div 20)} = 0.9605$. This means, for the

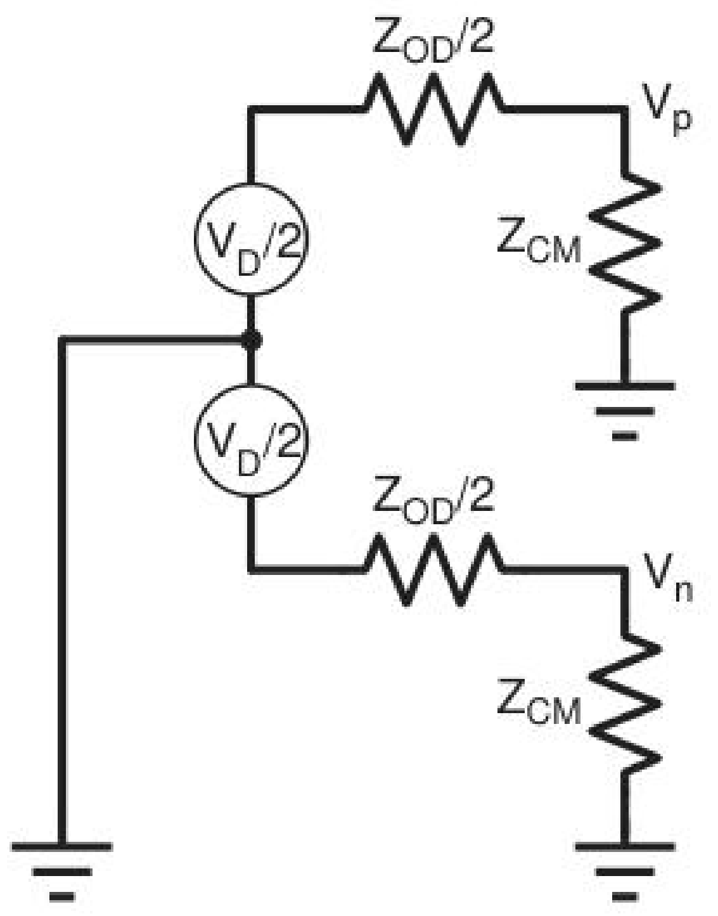

Fig.11: an equivalent circuit for balanced differential output (with no common-mode signal).

potential divider, $Z_{CM} \div (Z_{CM} + Z_{OD}/2) = 0.9605$, from which we find $Z_{OD} = 49\Omega$.

This is nearly exactly the THAT1646's maximum specified output impedance, in spite of the assumptions we made; a good result. The common-mode output impedance of the THAT1646 in parallel with 600Ω would reduce the effective value of Z_{CM} and lead to a lower estimated value of Z_{OD} if we took account of it.

Capacitive loads

In recent months, we discussed problems with driving capacitive loads using op amps in some detail. As well as its handling of differential signals, an advantage of the THAT1646 for driving the output of the frequency response measurement circuit is its ability to drive large capacitive loads without being destabilised.

Balanced audio cables commonly use a shielded twisted-pair (STP) structure. The two signal lines are carried on a twisted pair, which is shielded to reduce interference coupling to the signal wires. Using a twisted-pair structure ensures any interference coupling is equal on both wires and can be rejected by the line receiver.

These cables have both a conductor-to-conductor capacitance (C_{cc}) and conductor-to-shield capacitance (C_{cs}) – see Fig.12.

From the perspective of the differential output, the two C_{cs} capacitances are in series so have an effective capacitance of $C_{cs}/2$. The series shield capacitances are in parallel with the conductor-to-conductor capacitance. That means the total capacitance seen by the differential output is $C_{cc} + C_{cs}/2$. The capacitance values are usually specified per unit length (eg, in pF/m).

Fig.13 shows an LTspice schematic

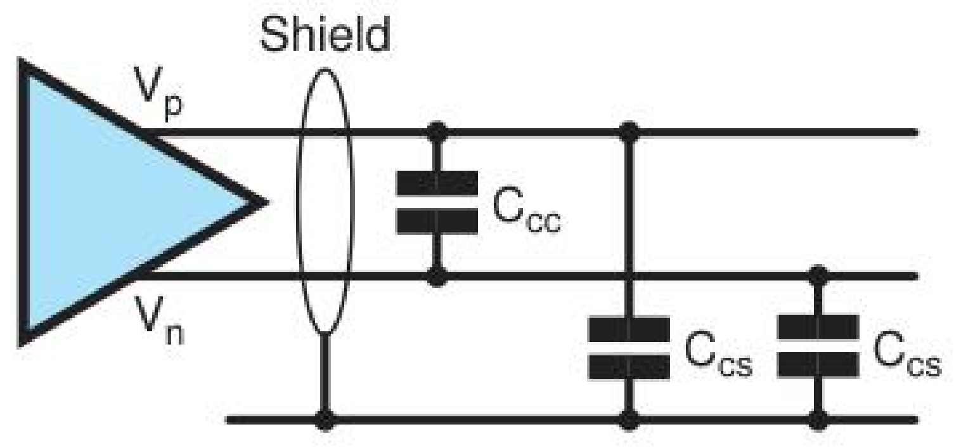

Fig.12: the capacitances on STP cable connected to a line driver.

that's similar to Fig.2 but with a capacitive load added, specifically a C_{cc} of 20nF and a C_{cs} of 20nF on each output. This is equivalent to an audio cable with a length in the low hundreds of meters, although of course part of the load capacitance will also come from any inputs to which the signal is connected.

This simulation was configured to drive the THAT1646 directly from the V1 source, which is producing fast-edged pulses. The OPAx134 produces a small amount of ringing with this input, so this avoids any influence from it. The op amp is kept in the circuit, as it is useful to also be able to rewire to see the behaviour of the circuit with a more realistic input to the line driver.

The results from the transient simulation are shown in Fig.14. The upper pane shows the input, the middle pane the two individual outputs, and the lower pane shows the differential output in a similar arrangement to Fig.3. Fig.15 provides zoomed-in detail of the step response and shows the input multiplied by two on the same pane as the output.

The results show that the THAT1646 is able to drive the capacitive load without any stability problems. Fig.15 shows that the response has no ringing and is slightly damped.

Fig.16 shows the AC analysis results on the same scale as Fig.4. The capacitive load reduces the bandwidth (to about 107kHz), as would be expected, but does not cause any significant change in the shape of the response. In particular, there is none of the gain peaking that we saw in recent articles concerning op amps that had difficulty driving capacitive loads.

The datasheet states that the THAT1646 is stable with any capacitive load, but some factors other than stability must be considered. As seen by comparing Fig.4 and Fig.16, capacitive loads reduce the bandwidth and the slew rate (output rate of change), which is related to the current drive capability. This needs to be taken into account.

Changing the voltage on a capacitor requires current – the faster the change or the larger the capacitance, the more current is required. If the current demand is too large, the THAT1646 will enter output current limiting and it will not be able to keep up with the waveform changes, resulting in distortion.

The THAT1646 datasheet provides more details on this with an example. Their example is similar to the situation in Fig.13 and is well within the THAT1646's capabilities.

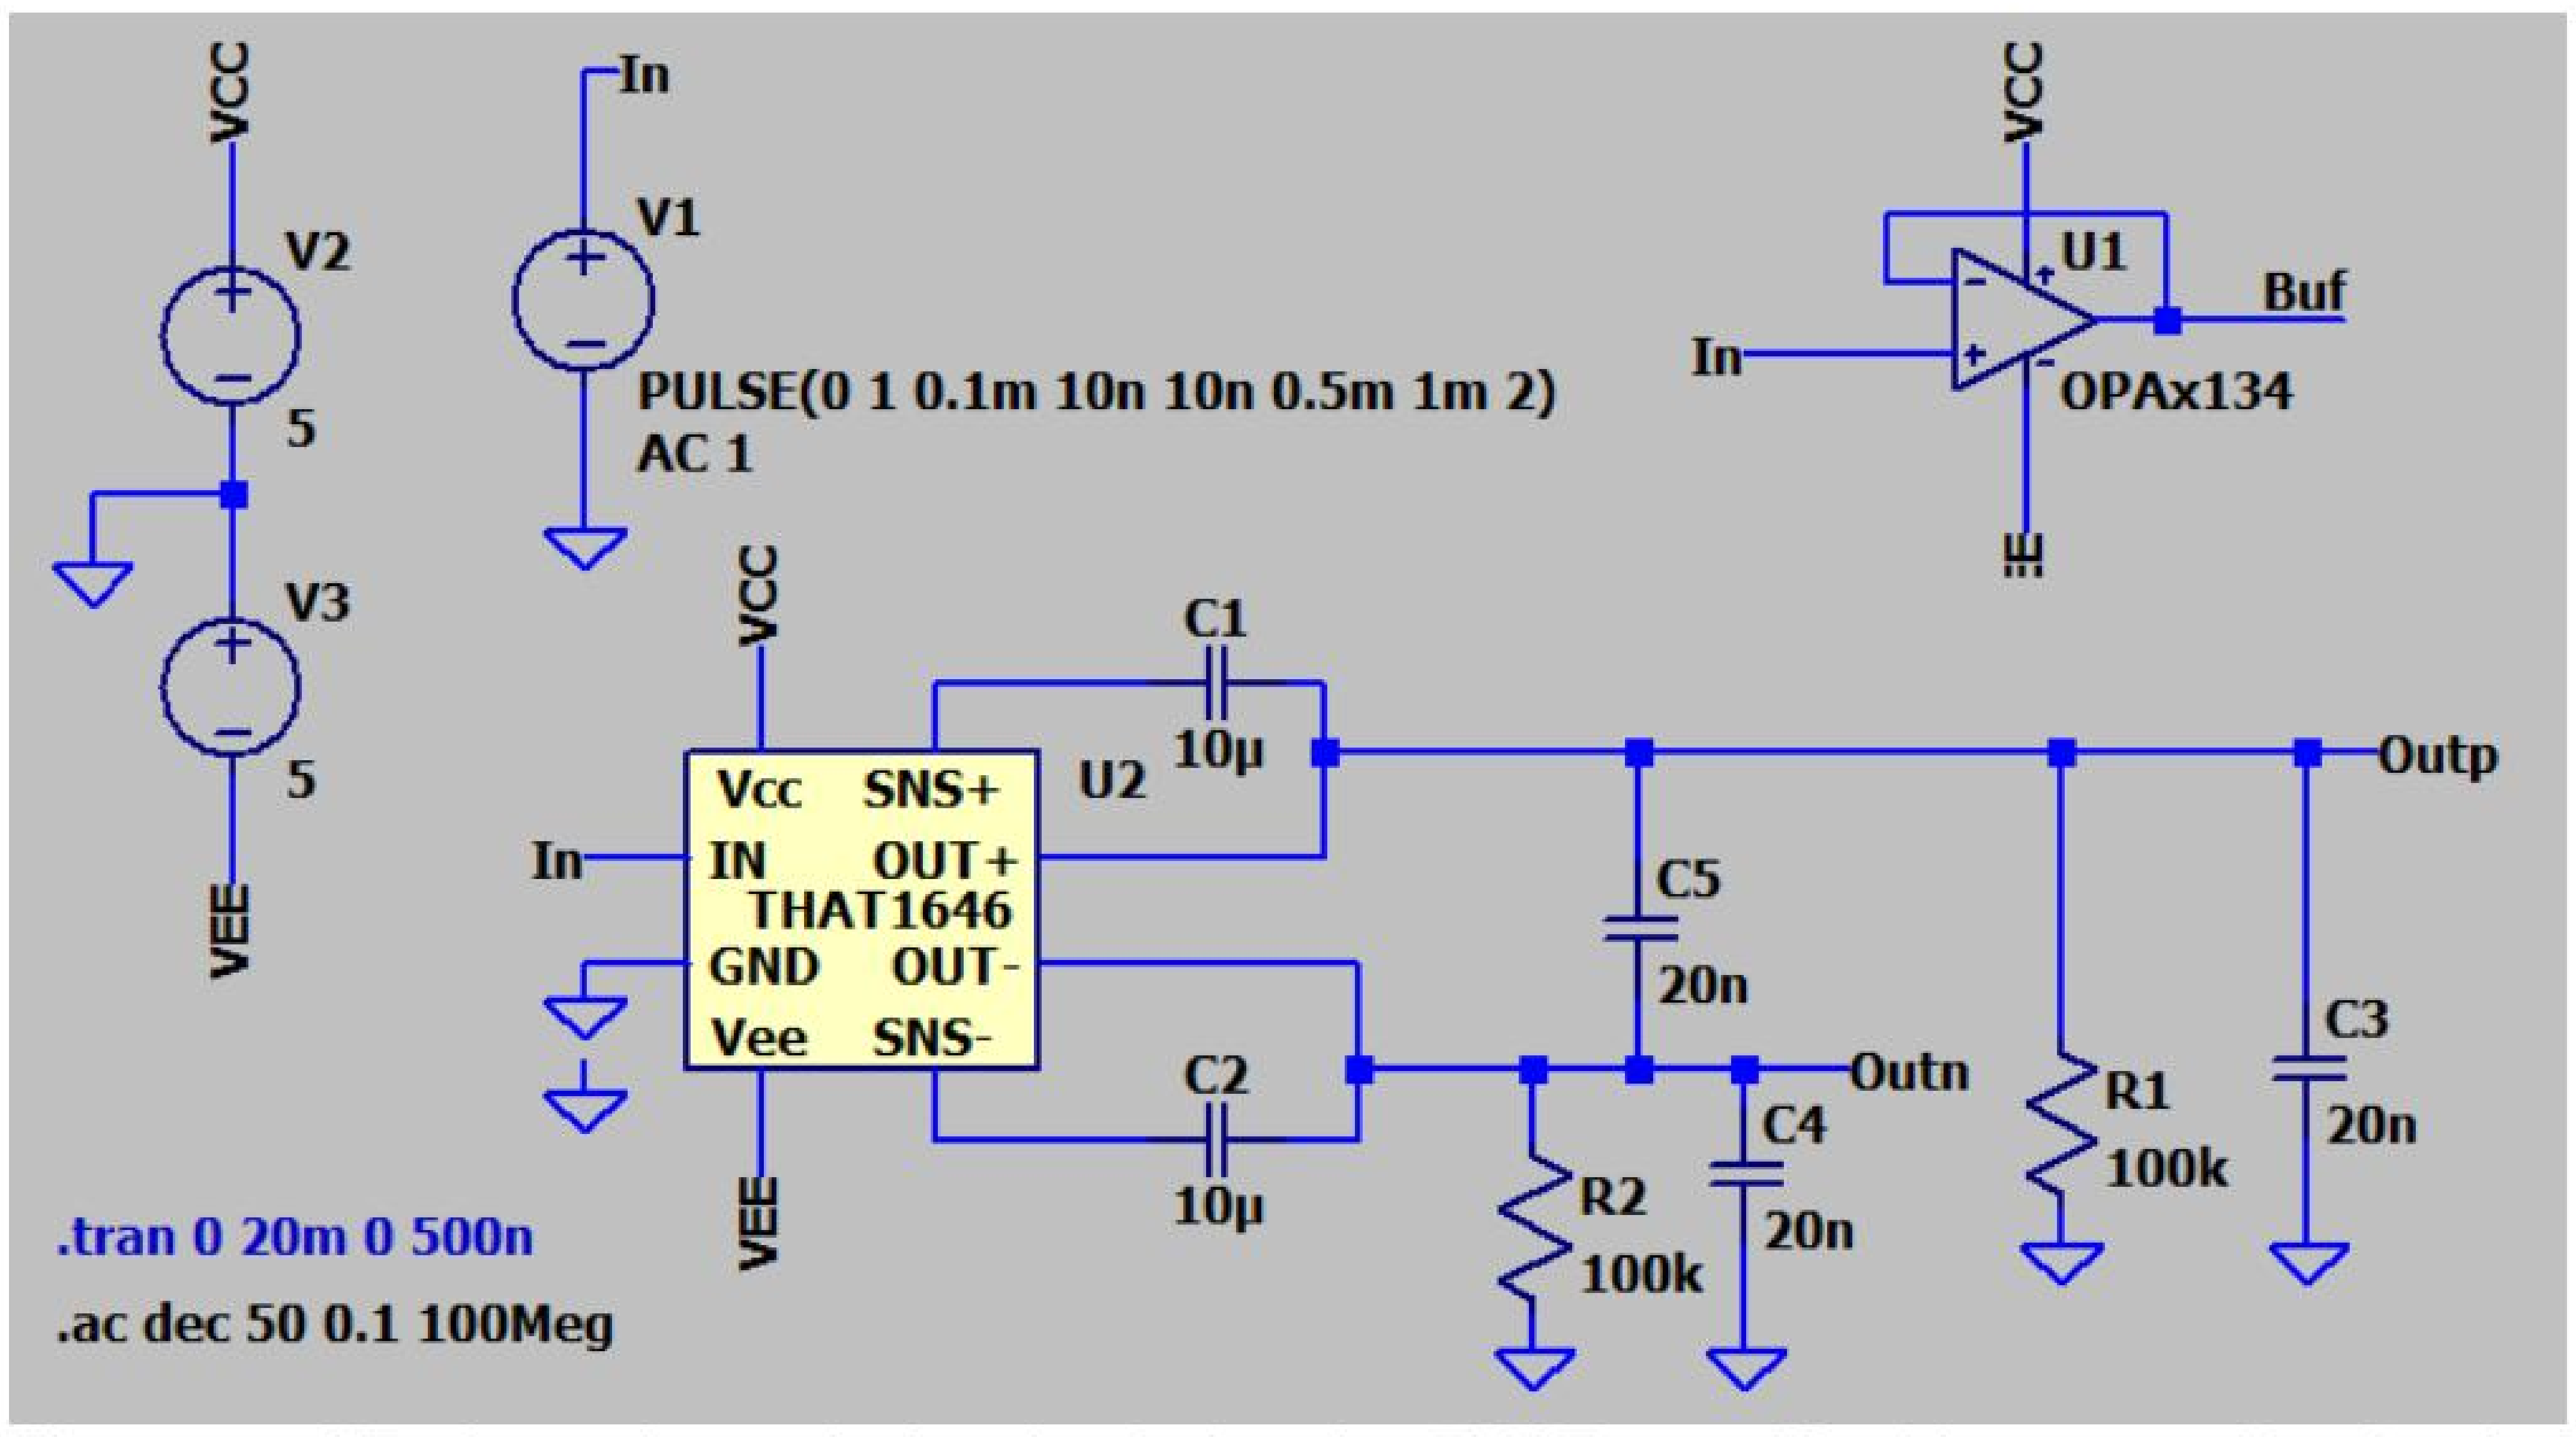

Fig.13: an LTspice schematic for simulating the THAT1646 IC with a capacitive load.

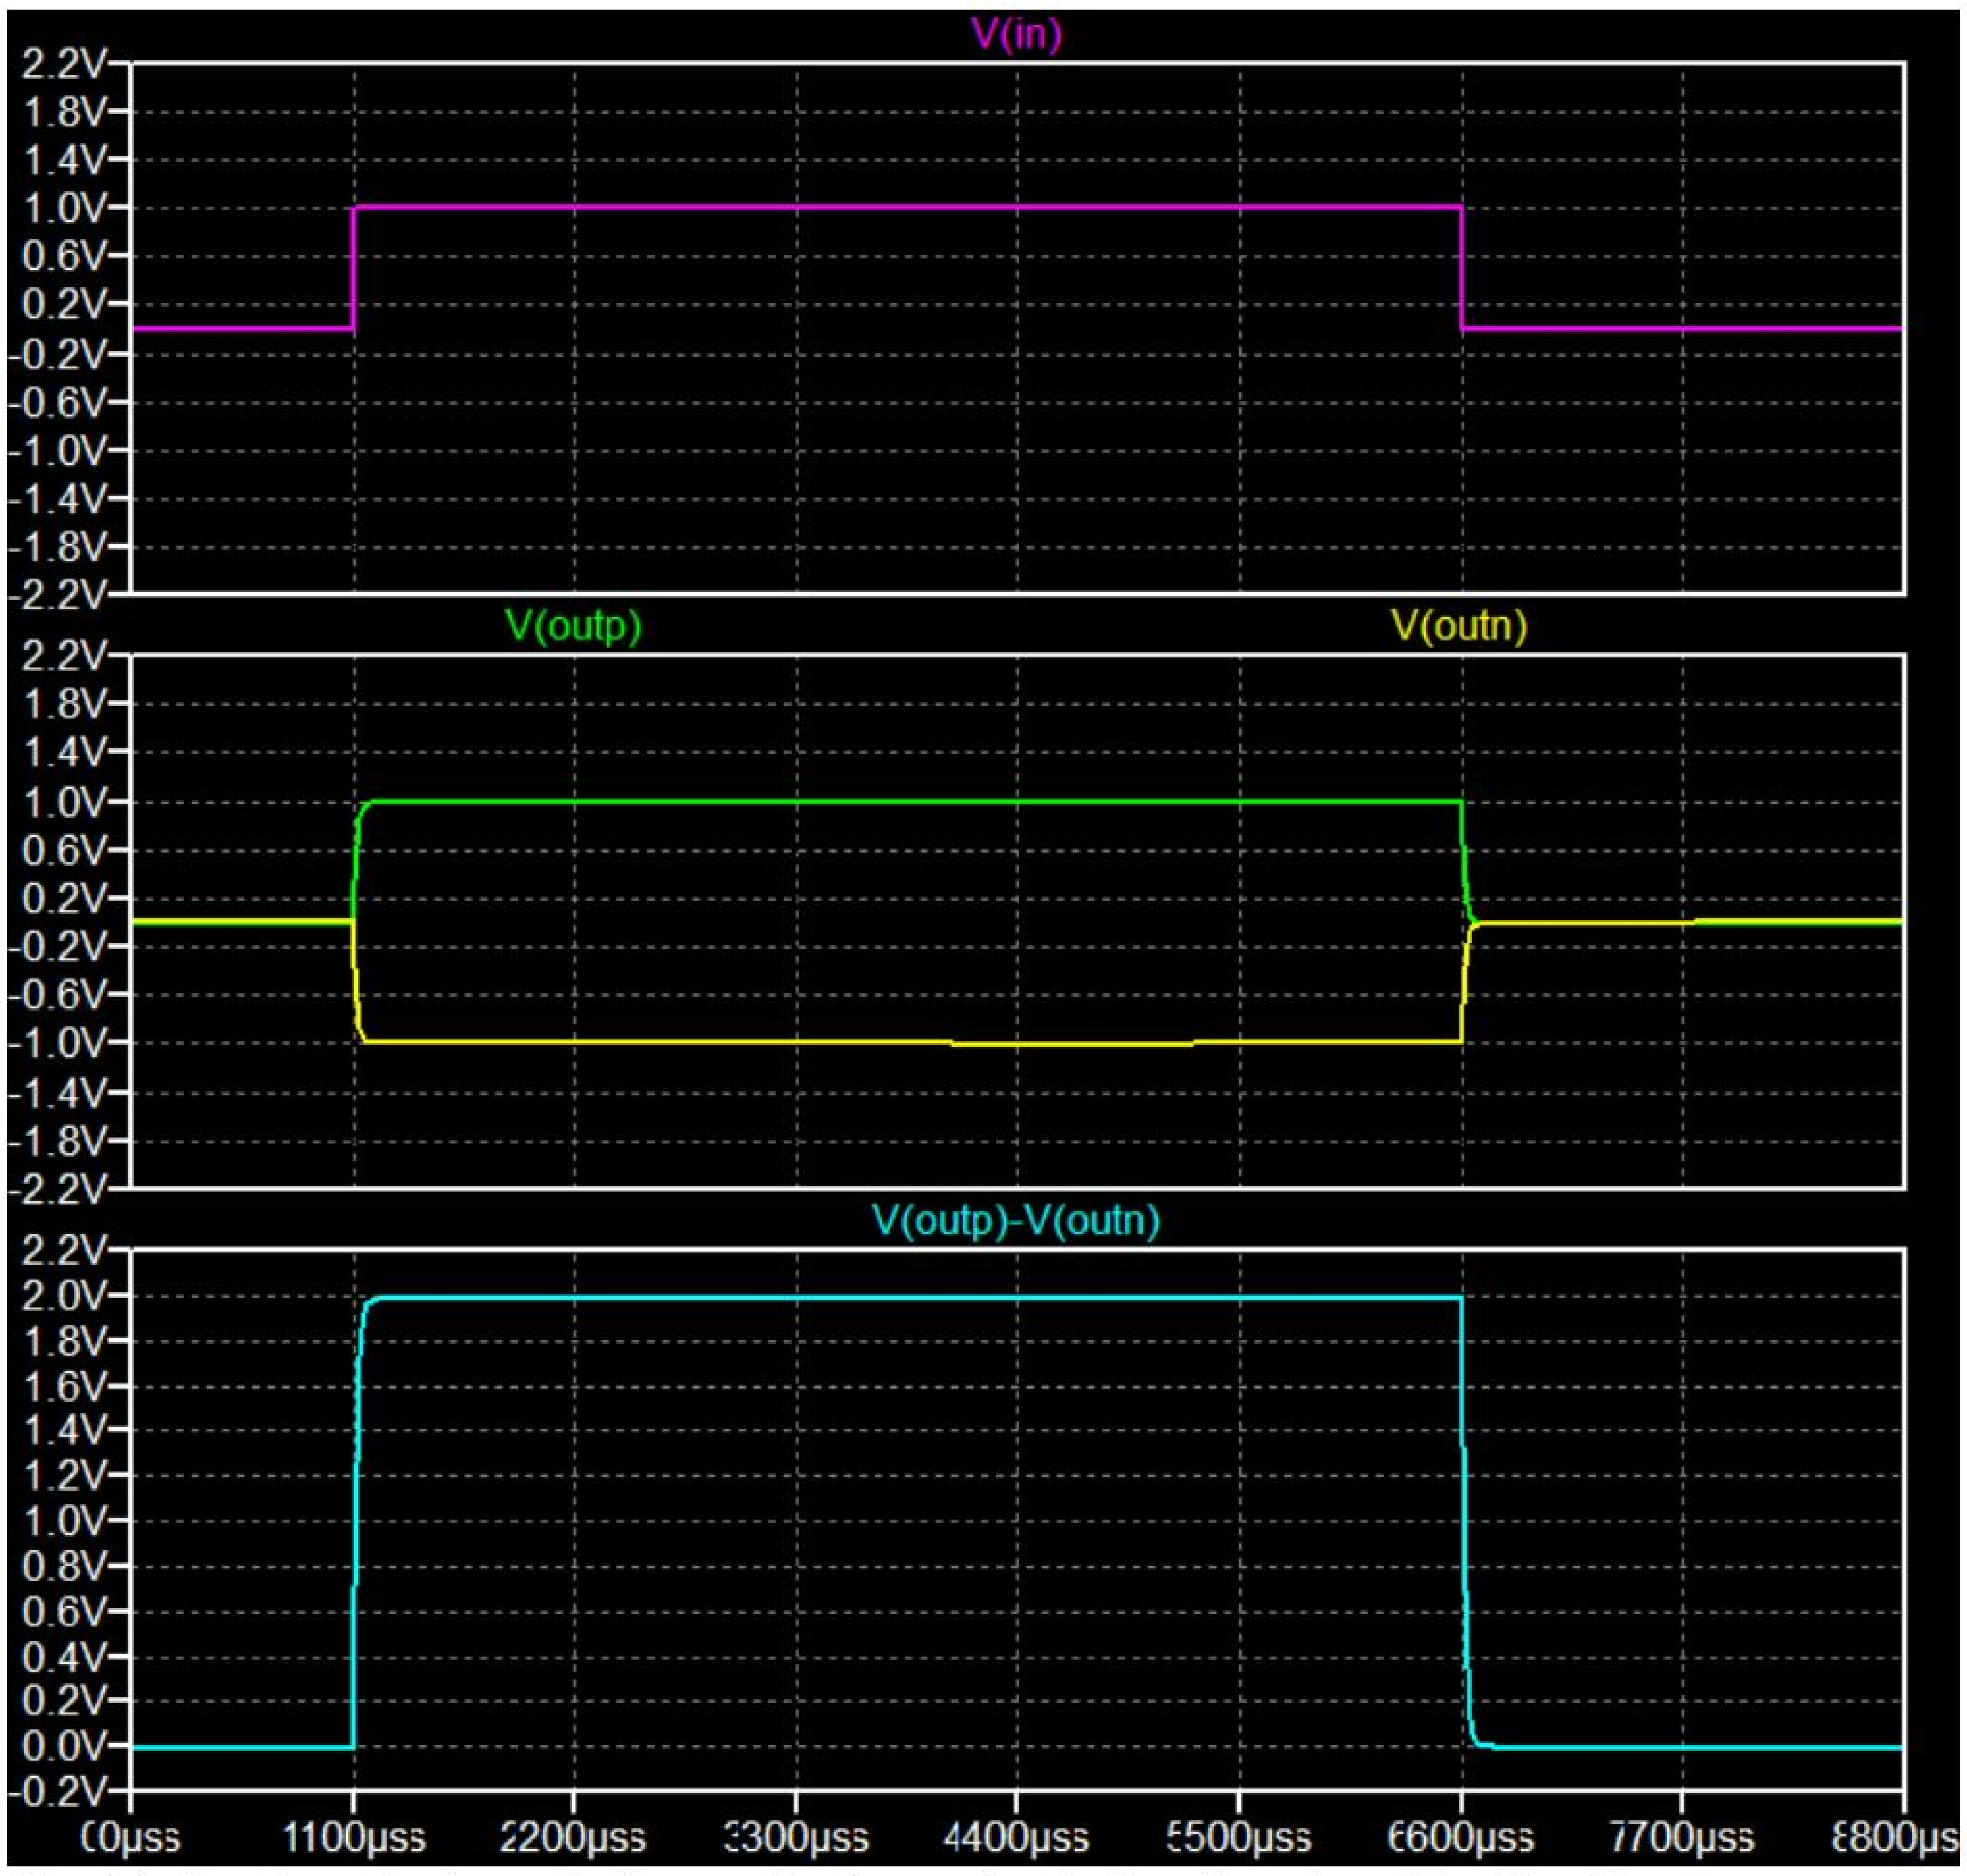

Fig.14: the transient analysis results from simulating the circuit in Fig.13.

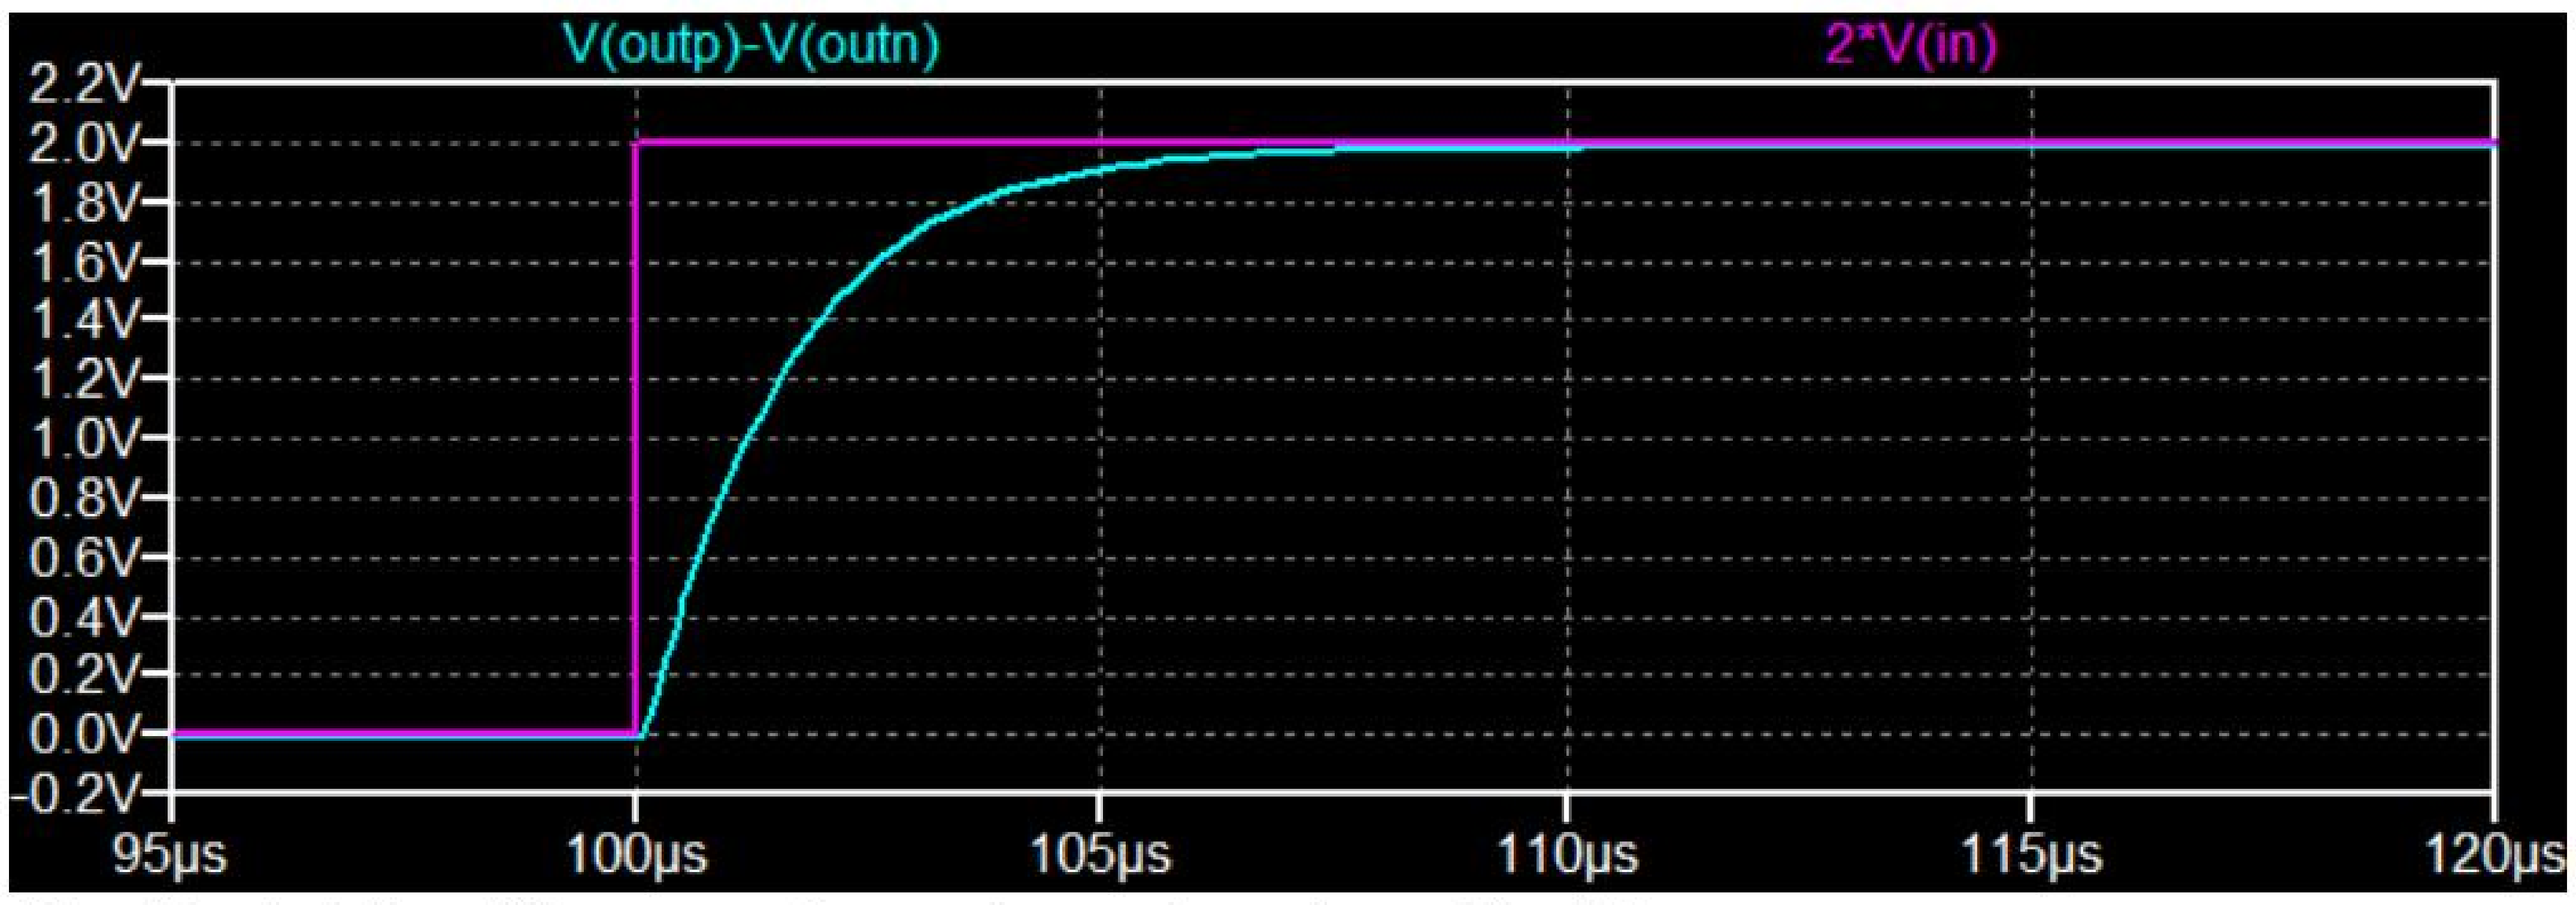

Fig.15: details of the waveform step edges from Fig.14.

Difference & differential amps

The THAT1646 contains a differential amplifier, and for a line driver and receiver combination to work well for a balanced line, we also need a good differential amplifier in the receiver. It is therefore useful to look in a bit more depth at the design and properties of these circuits.

Fig.17 shows a standard op amp based differential amplifier, which has a differential input and single-ended output. As with all op amp based amplifiers, negative feedback is applied to a very

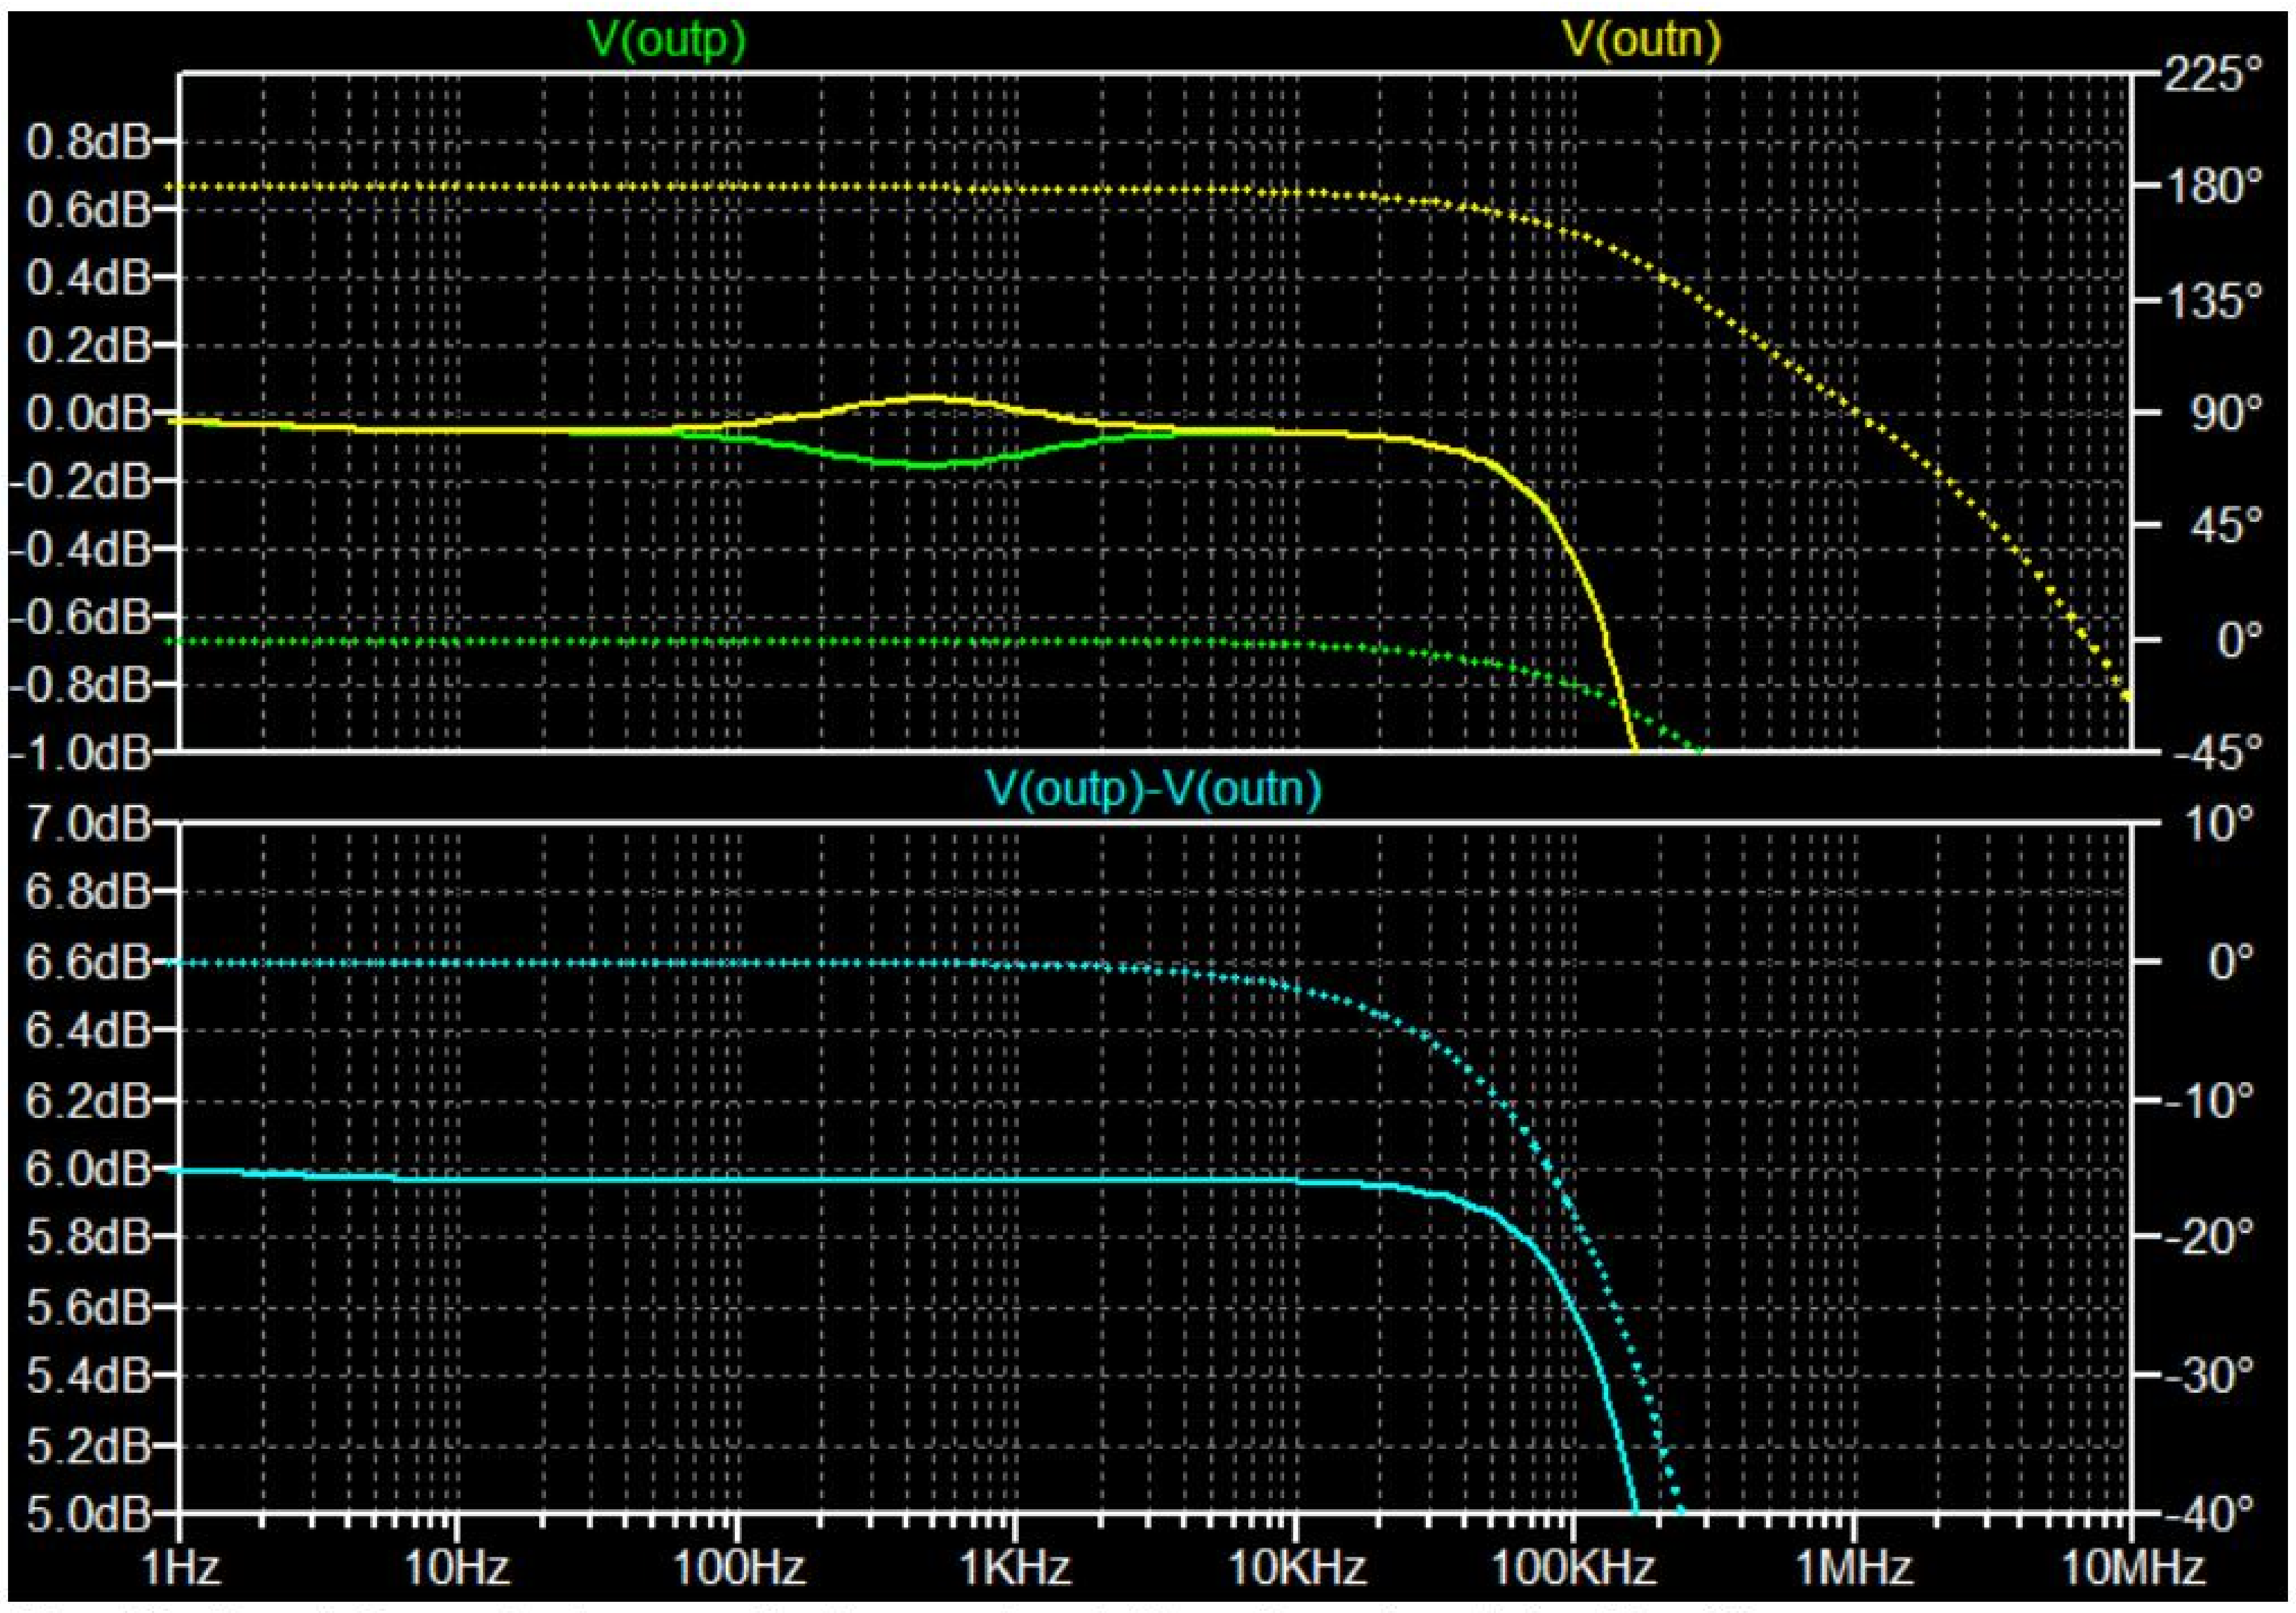

Fig.16: the AC analysis results from simulating the circuit in Fig.13.

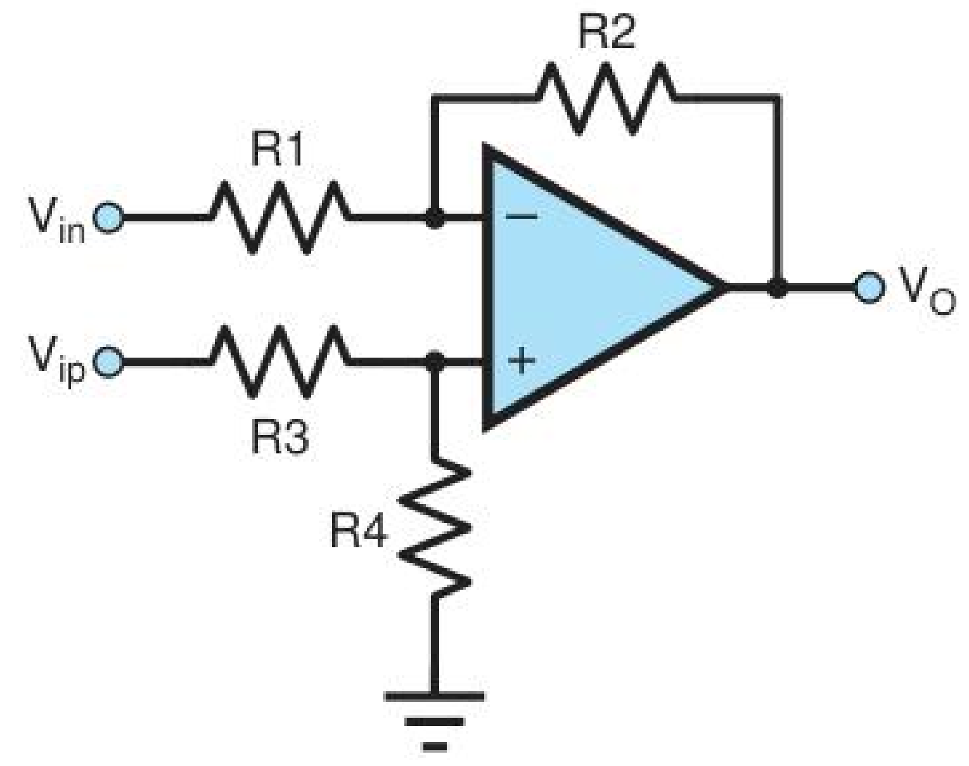

Fig.17 an op amp based differential amplifier (difference amplifier).

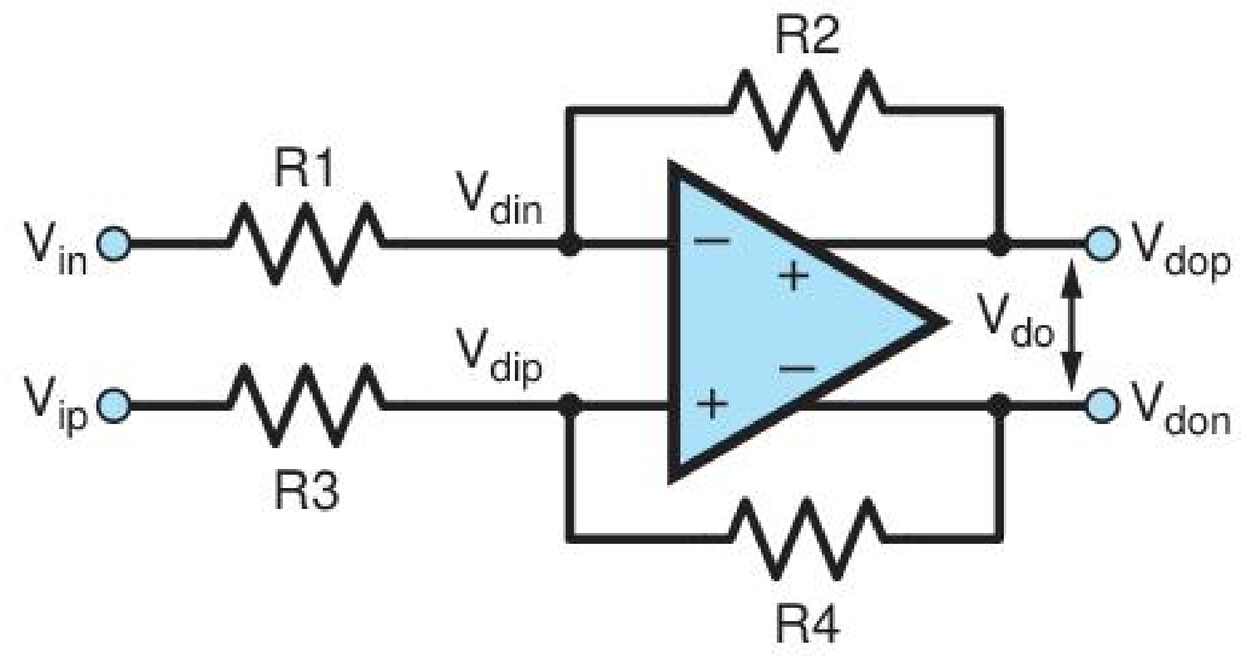

Fig.18: a fully differential amplifier.

high-gain op amp to set the gain of the circuit.

Fig.18 shows a fully differential amplifier comprising a fully differential op amp and two feedback networks. This is similar to a standard op amp circuit, except that the very high-gain amplifier has differential outputs as well as differential inputs.

As with standard op amp circuits, the feedback resistors set the gain. Negative feedback is applied via both inputs (from positive out to negative in, and negative out to positive in).

Both the circuits in Fig.17 and Fig.18 can be called differential amplifiers, but the terms difference amplifier (Fig.17) and fully differential amplifier (Fig.18) are used to distinguish them.

Difference amplifier operation

The difference amplifier in Fig.17 can be considered as a combination of a standard inverting and non-inverting amplifier, with resistors R_3 and R_4 acting as a potential divider on the input of the non-inverting amplifier (the standard circuit feeds the input directly to the op amp).

We can use a key piece of circuit theory called superposition to analyse the inverting and non-inverting aspects of the circuit separately.

For linear circuits, we can use superposition to deal with multiple sources – we set all sources except one to zero and find the effect of that source on its own (for example, its contribution to output voltage), then repeat with the other sources. Finally (because the circuit is linear), we can add up the individual contributions.

If we set $v_{ip} = 0$, the circuit in Fig.17 acts as an inverting amplifier with respect to v_{in}. Resistors R_3 and R_4 are in parallel, connecting the non-inverting input to ground, and do not affect the gain. The gain with respect to v_{in} is $-R_2 \div R_1$, in accordance with the standard inverting gain formula, so v_{in} makes a contribution to v_o of $-(R_2 \div R_1) \times v_{in}$.

If we set $v_{in} = 0$, the circuit in Fig.17 acts as a non-inverting amplifier, for which the gain is $1 + R_2 \div R_1$, using the standard gain formula. We can also write the gain as $(R_2 + R_1) \div R_1$.

The input signal (v_{ip}) is applied via the potential divider formed by R_3 and R_4, so to find the output, we multiply v_{ip} by the potential divider formula and the gain. Therefore, v_{ip} makes a contribution to v_{op} of $(R_4 \div (R_3 + R_4)) \times (1 + R_2 \div R_1) \times v_{ip}$.

The output from the circuit is the sum of these two contributions. If we choose $R_3 = R_1$ and $R_4 = R_2$ and substitute R_1 and R_2 for R_3 and R_4 in the v_{ip} contribution stated above, we get the overall output as:

$$v_o = \left(\frac{R_2}{R_1+R_2}\right)\left(\frac{R_1+R_2}{R_1}\right)v_{ip} - \left(\frac{R_2}{R_1}\right)v_{in}$$

The (R_2+R_1) terms cancel and the equation simplifies to:

$$v_o = \left(\frac{R_2}{R_1}\right)\left(v_{ip} - v_{in}\right)$$

The circuit amplifies the difference between its inputs by a gain of $R_2 \div R_1$.

Using equal resistors, as stated above, gives difference amplification, but equality of the resistors is not the only way to approach this. The use of ratios of component values can also be an important factor in circuit design and is significant here.

The circuit has two potential dividers, R_1 and R_2 for the feedback, and R_3 and R_4 for v_{ip}. If we define the two divider resistor ratios as $D_{12} = R_2 \div R_1$ and $D_{34} = R_4 \div R_3$, the contribution of v_{in} to the output is $D_{12}v_{in}$.

For v_{ip}, we first note that we can write the potential divider formula $R_4 \div (R_3+R_4)$ as $(R_4 \div R_3) \div (1 + R_4 \div R_3)$ by dividing both the numerator and denominator by R_3. This can be written as $D_{34} \div (1 + D_{34})$ using the divider ratio. We can then write the output voltage expression as:

$$v_o = \left(\frac{D_{34}}{1+D_{34}}\right)\left(1+D_{12}\right)v_{ip} - D_{12}v_{in}$$

Equal resistances of $R_3 = R_1$ and $R_4 = R_2$ mean that $D_{12} = D_{34}$, so $(1 + D_{12})$ cancels in the v_{ip} term, and the equation simplifies to:

$$v_o = D_{12}\left(v_{ip} - v_{in}\right)$$

We could also get this by substituting D_{12} into the expression for v_o but the full equation using D_{12} and D_{34} is useful because it shows us that the input-output relationship depends only on the resistor ratios, not on their absolute values. Equality of the ratios will provide correct operation even if the absolute values are not equal.

This is important because some fabrication techniques may be able to provide better ratio matching than it can control absolute values. We will discuss this more shortly.

The THAT1240 series

Although the focus of this article is line drivers, it is worth mentioning balanced line receivers. Examples include the INA134 and INA137 from Texas

Simulation files

Most *Circuit Surgery* columns feature the use of the free circuit simulation software LTSpice. It is used to support descriptions and analysis in *Circuit Surgery*.

The examples and files for this issue are available for download: https://pemag.au/link/acbx

Instruments, the SSM2141 and SSM2143 from Analog Devices and the THAT1240 series from THAT Corporation. These all comprise an op amp and four laser-trimmed resistors, as shown for the THAT1240 series in Fig.19.

The variants of these devices have different resistor values to obtain different gains; typically 0dB, ±3dB or ±6dB. The resistors on the ICs are not fully connected to form a difference amplifier, as in Fig.17, so that they can be wired externally to achieve different gains (eg, +3dB or -3dB).

Fully differential amplifiers

Referring to Fig.18, we have a fully differential amplifier which has a (very large) differential gain of A_D, so $v_{do} = (v_{dop} - v_{don}) = A_D(v_{dip} - v_{din})$. Here, v_{dip} and v_{din} are the voltages directly at the fully differential op amp's inputs, but we need to find the differential output in terms of the circuit input voltages, v_{ip} and v_{in}.

The op amp's input voltages depend on the relevant output and circuit input voltages connected via the resistors. The two resistor pairs form potential dividers that act the opposite way round with respect to the circuit's inputs and the op amp's outputs. We can use superposition again to sum the two contributions at each op amp input.

For v_{din} with respect to v_{dop}, R_1 and R_2 form a potential divider in which R_1 is the lower resistor, so the contribution of v_{dop} to v_{din} is found by applying the potential divider formula to get $(R_1 \div (R_1 + R_2)) \times v_{dop}$. We can define F_{xy} as the 'potential divider fraction' for resistors x and y, with x as the lower resistor, so that $F_{12} = R_1 \div (R_1 + R_2)$. This will help keep our equations tidy.

Using this approach and summing the two contributions, we find $v_{din} = F_{12} \times v_{dop} + F_{21} \times v_{in}$ and similarly, $v_{dip} = F_{34} \times v_{don} + F_{43} \times v_{ip}$. Like the difference amplifier, we use equal resistor ratios, so $F_{12} = F_{43}$ and $F_{21} = F_{34}$. That means $v_{dip} = F_{12} \times v_{don} + F_{21} \times v_{ip}$. We can now substitute these expressions for v_{dip} and v_{din} in the equation for v_{do} above to get:

$$v_{do} = A_D\left[\left(F_{12} v_{don} + F_{21} v_{ip}\right) - \left(F_{12} v_{dop} + F_{21} v_{in}\right)\right]$$

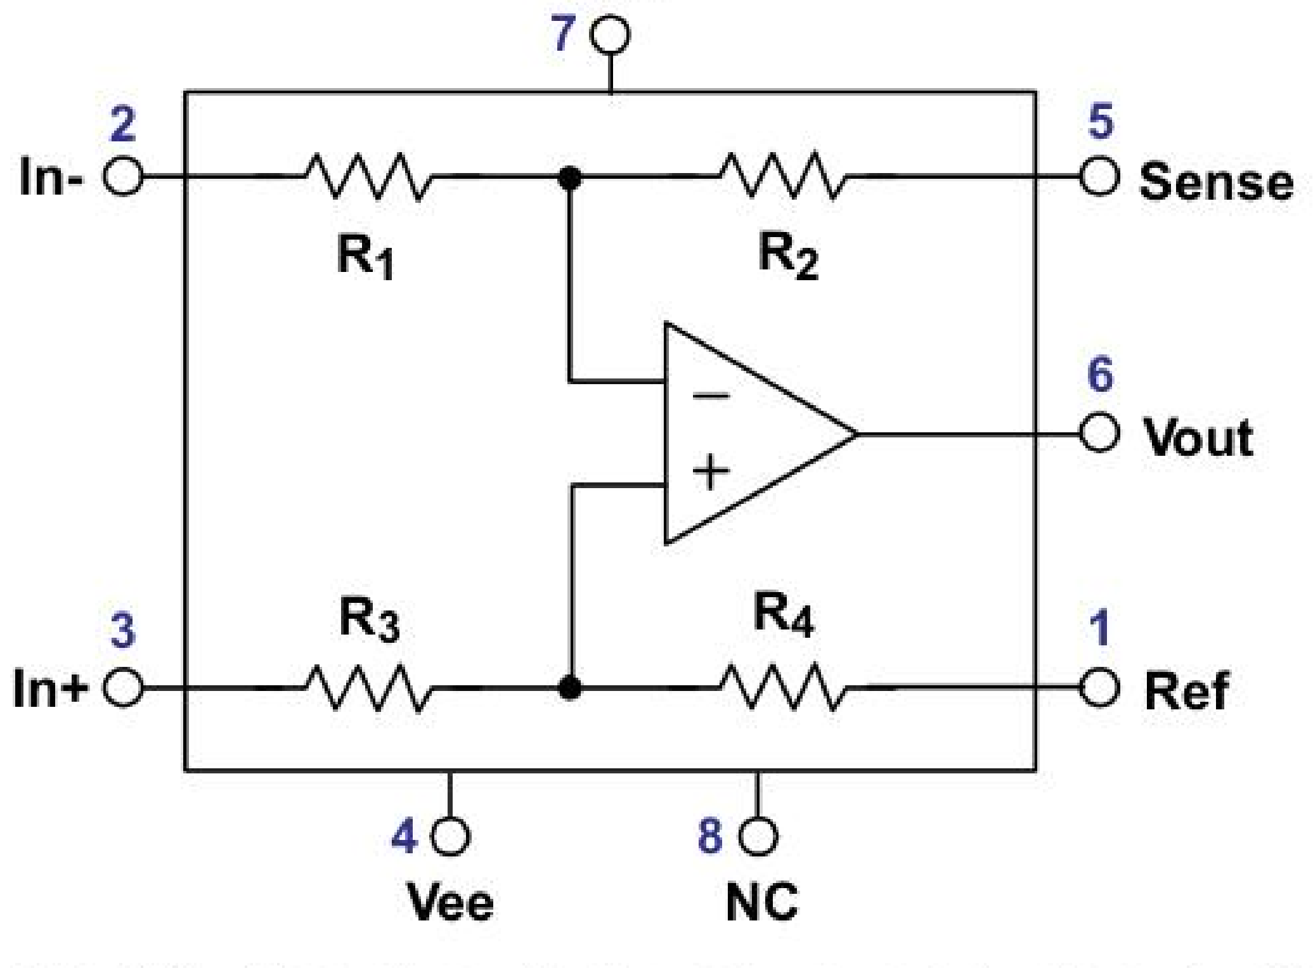

Fig.19: the circuit structure and pinout of the THAT1240 IC (from the data sheet).

Rearranging to get the inputs and outputs together, then replacing the input difference with v_{do} gives:

$$v_{do} = A_D\left[\left(F_{12} v_{don} - F_{12} v_{dop}\right) + \left(F_{21} v_{ip} - F_{21} v_{in}\right)\right]$$

$$v_{do} = A_D\left[-F_{12} v_{do} + F_{21}\left(v_{ip} - v_{in}\right)\right]$$

Dividing by A_D and moving both output terms to the left-hand side results in:

$$\frac{v_{do}}{A_D} + F_{12} v_{do} = F_{21}\left(v_{ip} - v_{in}\right)$$

Given that A_D is very large (tens of thousands, or much more) and the F_{12} fraction is typically no more than tenths or hundredths, and often larger, we can assume that the $v_{do} \div A_D$ term is much smaller than $F_{12}v_{do}$ and can be ignored. We then divide by F_{12} to get:

$$v_{do} = \frac{F_{21}}{F_{12}}\left(v_{ip} - v_{in}\right)$$

If we divide two opposite-direction potential divider formulae, the resistor sums cancel, and we are left with the ratio of the two individual resistors, so:

$$v_{do} = \frac{R_2}{R_1}\left(v_{ip} - v_{in}\right)$$

Note the similarity with the standard inverting amplifier, which has the same feedback-over-input resistor formula for its gain. For the THAT1646, R_1 = 5kΩ and R_2 = 10kΩ (see Fig.1), so $R_2 \div R_1 = 2$. Also, v_{in} is grounded (equal to zero), so $v_{do} = 2v_{ip}$, that is, a gain of 2× (+6dB) with respect to the differential output, as seen in the simulations above.

Common-mode rejection

The common-mode rejection ratio (CMRR) is defined as the ratio of differential (A_D) to common-mode gain (A_{CM}), and is often expressed in decibels, calculated as $20\log_{10}(A_D \div A_{CM})$.

Ideally, differential amplifies have infinite CMRR. Unfortunately, both the difference amplifier and differential amplifier circuits have a CMRR that is highly dependent on the gain-setting resistors' tolerance, even when using ideal op amps having infinite CMRR.

Resistor tolerance (T) for a resistor of value R is defined as $T = \Delta R \div R$, where ΔR is the largest deviation from R (positive or negative) for which the resistor is considered as being in specification. Usually, this is multiplied by 100 to give a percentage value when tolerances are stated, but analysis and resulting equations will normally use the absolute value of T.

This may be stated without units, or with units of Ω/Ω. The worst-case actual or absolute value of an individual resistor (R_A) is given by $R_A = R_N(1 \pm T)$, where R_N is the nominal resistor value.

When analysing the effect of component tolerance on circuit characteristics, we can use the worst-case combinations to find the maximum impact. For example, consider a basic inverting amplifier with gain $-R_2 \div R_1$ using nominal values of 10kΩ and 1kΩ to give a gain of -10.

The worst-case deviation in gain due to resistor tolerance would occur if the two resistors differed from nominal by the full tolerance in opposite directions. For example, for 1% resistors (T=0.01), we get -(10kΩ×1.01)÷(1kΩ×0.99) = -10.202 or -(10kΩ×0.99)÷(1kΩ×1.01) = 9.802. This is about a 2% error in both cases.

Analysis of the effect of resistor tolerance on the CMRR of both difference and differential op amp amplifiers is more complex than for the gain of the amplifier we have just looked at. Although the final formula is straightforward, the analysis leads to some large equations, so we will not go through the details here.

You can find the details of this, and related analyses, in two documents: Design Note 1023 from Linear Technology, now Analog Devices (**www.analog.com/media/en/reference-design-documentation/design-notes/dn1023f.pdf**) and Application Note SBOA582 from Texas Instruments (**www.ti.com/lit/SBOA582**).

For difference and differential amplifiers with a circuit gain of G built using four independent resistors of tolerance T, the worst-case CMRR due to the resistors (CMRR_R) is $(G + 1) \div 4T$.

Some example CMRR values for various amplifier gains and resistor tolerances are shown in Table.1. General-purpose op amps typically have CMRR values in the range of 70-100dB but higher-precision devices may range up to around 130dB.

Given that CMRR is important for balanced line receivers, we would not want to undermine it through poor resistor choices. Table.1 implies that we would need resistors with tolerances of ±0.01% or better to prevent the resistor-based CMRR from dominating the overall CMRR.

Resistors with ±0.01% and ±0.001% tolerance are available, but they cost significantly more than standard resistors. Even if cost is not a concern, using individual discrete resistors is not the best option.

Table 1: differential amplifier CMRR due to resistor tolerance for various gains.

Resistor tolerance	Gain		
	1×	10×	20×
5%	20dB	35dB	40dB
1%	34dB	49dB	54dB
0.1%	54dB	69dB	74dB
0.01%	74dB	89dB	94dB
0.001%	94dB	109dB	114dB

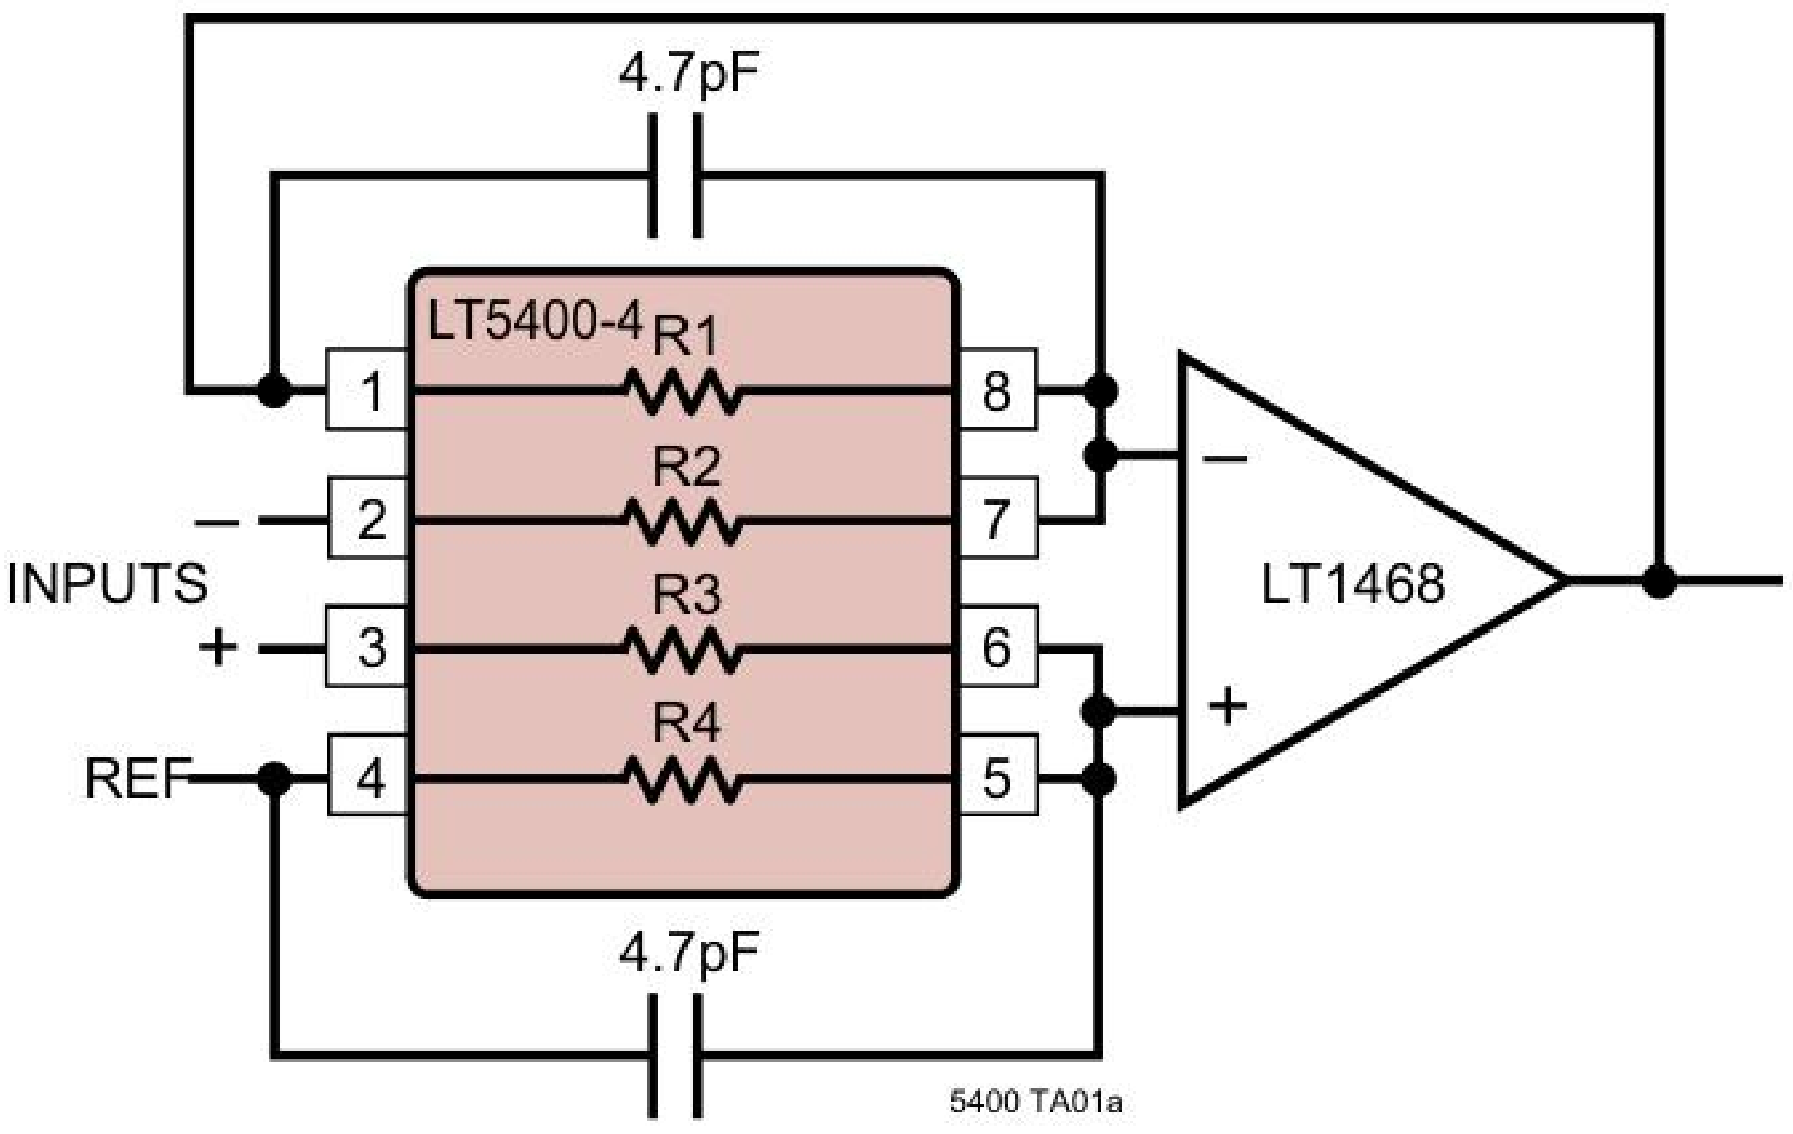

Fig.20: a difference amplifier using an LT5400 resistor network (from the data sheet).

Of course, if the resistors are well matched, their ratios will be well matched too, and we will have accurate gain and good CMRR. However, even if the resistors are perfectly matched at a given time, they may not stay that way. If they are subject to different temperatures, their values will shift differently. Also, they may age differently.

Effects such as this can degrade the CMRR over time or in changing environmental conditions.

Ratio matching

As discussed above, it is the resistor ratios that really matter, not the absolute values. Fortunately, it is possible to maintain a resistor ratio even if absolute values change (eg, with temperature).

This can be done to some extent with discrete resistors on a PCB by mounting them close together and avoiding temperature gradients. However, close placement is more effective for integrated devices, which also tend to have closer matching of temperature coefficients. IC designers also have access to layout techniques that help provide very good matching and ratio stability.

Another advantage of integrating the resistors in a differential amplifier is that on-chip resistors can be adjusted during manufacture to obtain very precise ratios. This approach uses thin-film resistors formed from material deposited on the chip. The resistances can be adjusted by burning away some of the material using a laser. This increases the resistance.

The process can be performed while the device is being tested on the wafer before the individual chips are cut from it.

The THAT1606 and THAT1646 line drivers and the various balanced line receiver ICs mentioned above all use laser-trimmed thin-film resistors to help ensure good performance.

One disadvantage of on-chip resistors is that a limited number of gain values and circuit configurations are available. Using op amps with discrete resistors has more flexibility, but as just discussed, will limit performance.

An alternative is to use a precision-matched resistor network – integrated devices that contain a number of resistors with accurate ratios. Like high-precision individual resistors, these devices are relatively expensive.

An example is the LT5400 quad matched resistor network from Analog Devices, which is available in eight variants with different resistor values and five ratio options. An example difference amplifier circuit using the LT5400 is shown in Fig.20.

The A-grade version of the LT5400 has ±0.01% matching of resistors with a resistor matching ratio temperature drift of 0.2 parts per million per degree Celsius (0.2ppm/°C). The absolute value tolerance is only ±7.5%, but as the previous discussion has indicated, it is the ratio that matters in differential amplifiers.

Models of the LT5400 are available in the Misc section of the LTspice component library.

A better THAT1646 circuit

The circuit in Fig.2 is a basic version of the THAT1646 line driver circuit.

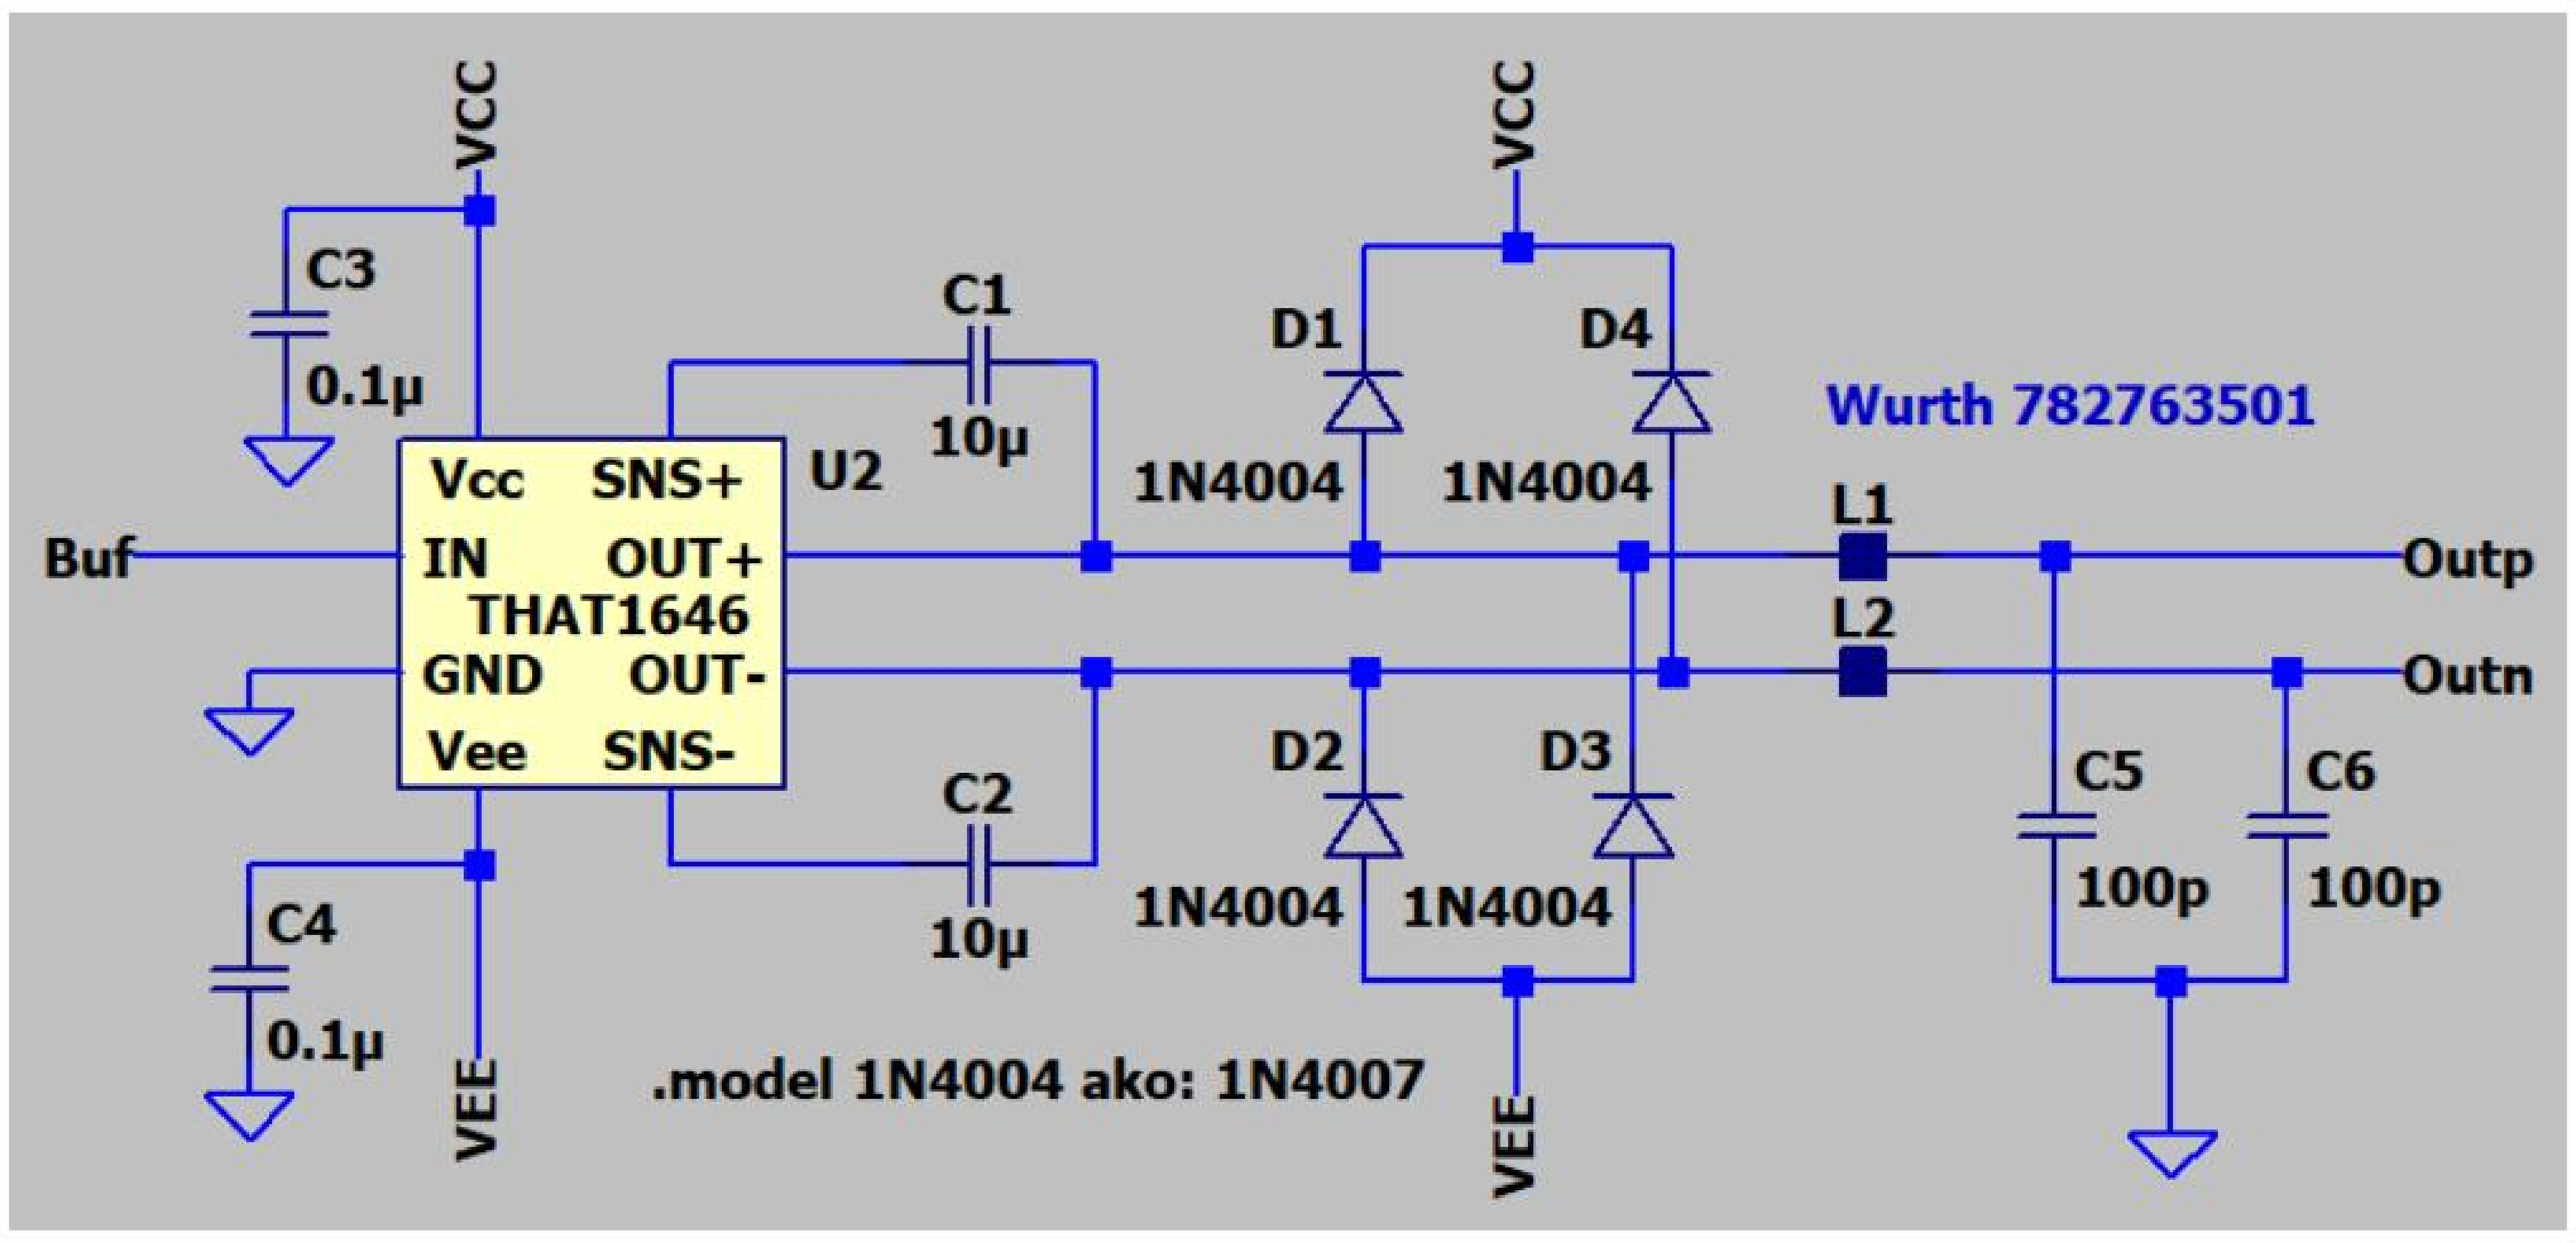

Fig.21: a THAT1646 based circuit with output and RFI protection.

Fig.21 has some additional components to provide output and radio frequency interference (RFI) protection. The additions are recommended in the datasheet.

The circuit in Fig.21 can be simulated with the addition of the supplies, sources and buffer from Fig.2. Decoupling capacitors **C3** and **C4** are required in a real circuit, but will have little effect in the simulation, due to the ideal voltage sources supplying the circuit.

Inductors **L1** & **L2** and capacitors **C5** & **C6** form two filters to prevent any RFI picked up by a cable on the output from getting into the circuit.

The 782763501 is an SMD (M3216/1206) ferrite bead from Würth Elektronik, which is available from the "Select Ferrite Bead" menu via right-clicking the schematic symbol. It has a 500Ω impedance at 100MHz. This specific choice is not critical; other ferrite beads would also be suitable.

Diodes **D1** to **D4** provide protection against the possibility that the circuit could be connected to a balanced input providing 48V phantom power. The datasheet circuit uses 1N4004 diodes, but LTspice only has the similar (higher voltage rating) 1N4007 devices in the library. The schematic is labelled using 1N4004 but references the 1N4007 model via the ako: (a kind of) model construct.

Phantom power circuits are used with condenser microphones that have built-in preamps. The 48V power is connected to both signal wires via a couple of resistors – thus appearing as a common-mode signal on the balanced line, and can be obtained from a centre-tapped secondary of the transformer in the microphone.

With 48V active, the line receiver input capacitors can discharge via the THAT1646 and may cause damage. As 48V is greater than the supply voltage, the diodes will conduct and absorb the energy instead.

Conclusion

By combining a balanced line driver like the THAT1646with a difference amplifier like that shown in Fig.21, we obtain a low-cost, robust system capable of sending audio over long cables with little degradation or EMI/noise pickup.

We do have to be careful with the selection of resistors in our difference amplifier in the receiver to keep the CMRR high. Otherwise, the ability of the system to reject noise pickup will be degraded. Luckily, numerous monolithic ICs and resistor networks are available to allow us to do this, as using individual resistors for this job is far from ideal.

Tracks and tribulations

Techno Talk

A ramble through the evolution of PCB layout, from tape and scalpels to overcaffeinated autorouters and today’s AI-powered tools that are starting to earn their keep.

Max the Magnificent

There was once a time (“gather around, children”) when designing a printed circuit board (PCB) didn’t involve clicking a mouse, dragging traces across a screen, or pressing a button marked “Autoroute and Pray.” In those far-off days of yore, PCB layout was a hands-on, eyes-straining, finger-sticking affair performed on large sheets of translucent drafting film.

Instead of pixels, we used precisely manufactured rolls of adhesive tape that came in various widths, each corresponding to a different track size. Want a 10-mil trace? Reach for the 10-mil tape. (In this case, “mil” means one thousandth of an inch. The term comes from “milli-inch,” not millimetre, which has caused no end of confusion over the years.)

Pads and component footprints were not dragged from a library—they were rub-down transfers. You’d carefully position them over the layout and burnish them into place using a stylus, all the while muttering arcane incantations to the gods in the hope that they landed exactly where intended, because lifting and repositioning them was… let’s say “suboptimal”.

Working at 2× or 4× scale was common practice, which meant your “small” board might occupy something the size of a dining table. Once complete, the artwork would be photographically reduced to its final dimensions, at which point any tiny misalignment you hadn’t previously noticed would suddenly become glaringly obvious.

The most common tape colour was black, but for those working with double-sided boards (gasps of astonishment), an alternative was to use red tape for one side (usually the top) and blue tape for the other (often the bottom). This made it much easier to visually separate the layers and avoid going cross-eyed when your sheets of drafting film were stacked on a light table.

On the one hand, this was a painful way of doing things. People joked that you never saw a layout designer wearing a woolly pullover, even in the coldest weather. This was because you didn’t want any errant fibres to become attached to your film, only to be immortalised as almost invisibly thin “wires” destined to cause hard-to-diagnose faults by shorting out real tracks.

On the other hand, there was something rather wonderful about it all. You could see the board taking shape beneath your hands. Every track was placed with intent. Every component had a physical presence. And it was as much an artisanal craft as it was an engineering discipline.

System overheating!

I remember circa 1982 when the company I was working for invested in a proper computer-aided design (CAD) system for PCB layout. This involved two large cathode-ray tube (CRT) screens mounted on top of (and integrated into) a meaty metal table, which also included a chunky workstation.

The screen on the left was black-and-white; the one on the right was in colour. We used a keyboard and command-line interface (CLI) associated with the left-hand screen to issue cryptic commands like V (for “via”) 1250,875 or T (for “track”) 1000,500 1500,500 to instruct the system to place a via at a specific X-Y location or to draw a track segment between two points.

As each command was executed, the corresponding object appeared on the right-hand screen.

This system was not as user-friendly as one might have hoped, especially given how much of our hard-earned lucre we’d invested in it. One day, the right-hand screen started flashing an enigmatic message: SE392. The text was in red on a white background, so we knew it had to be something important.

The system had been delivered with a veritable cornucopia of documentation. We’re talking about an entire shelf groaning under the weight of thick, ring-bound manuals. We finally located the “Messages” section in the first manual, which helpfully outlined the different categories. This was where we discovered that “SE” indicated a “System Error”.

It also directed us to an appendix in the last manual that contained the messages themselves. There were hundreds of them. When we finally tracked down SE392, we discovered that its corresponding description tersely read: “System Overheating”.

It was right. When we turned back to look at the workstation, we discovered it was on fire!

Push and shove

As the years went by, the computer-aided layout tools grew increasingly advanced. One great advancement was the use of interactive graphics. Suddenly, with the aid of a mouse (or, in some cases, a light pen or digitising tablet), we could place components and route tracks directly on the screen, which felt nothing short of magical.

Even so, PCB layout was still somewhat constrained. Routing was largely “Manhattan” in style, with tracks running horizontally or vertically and meeting at crisp 90° corners. The idea of casually introducing 45° segments came later

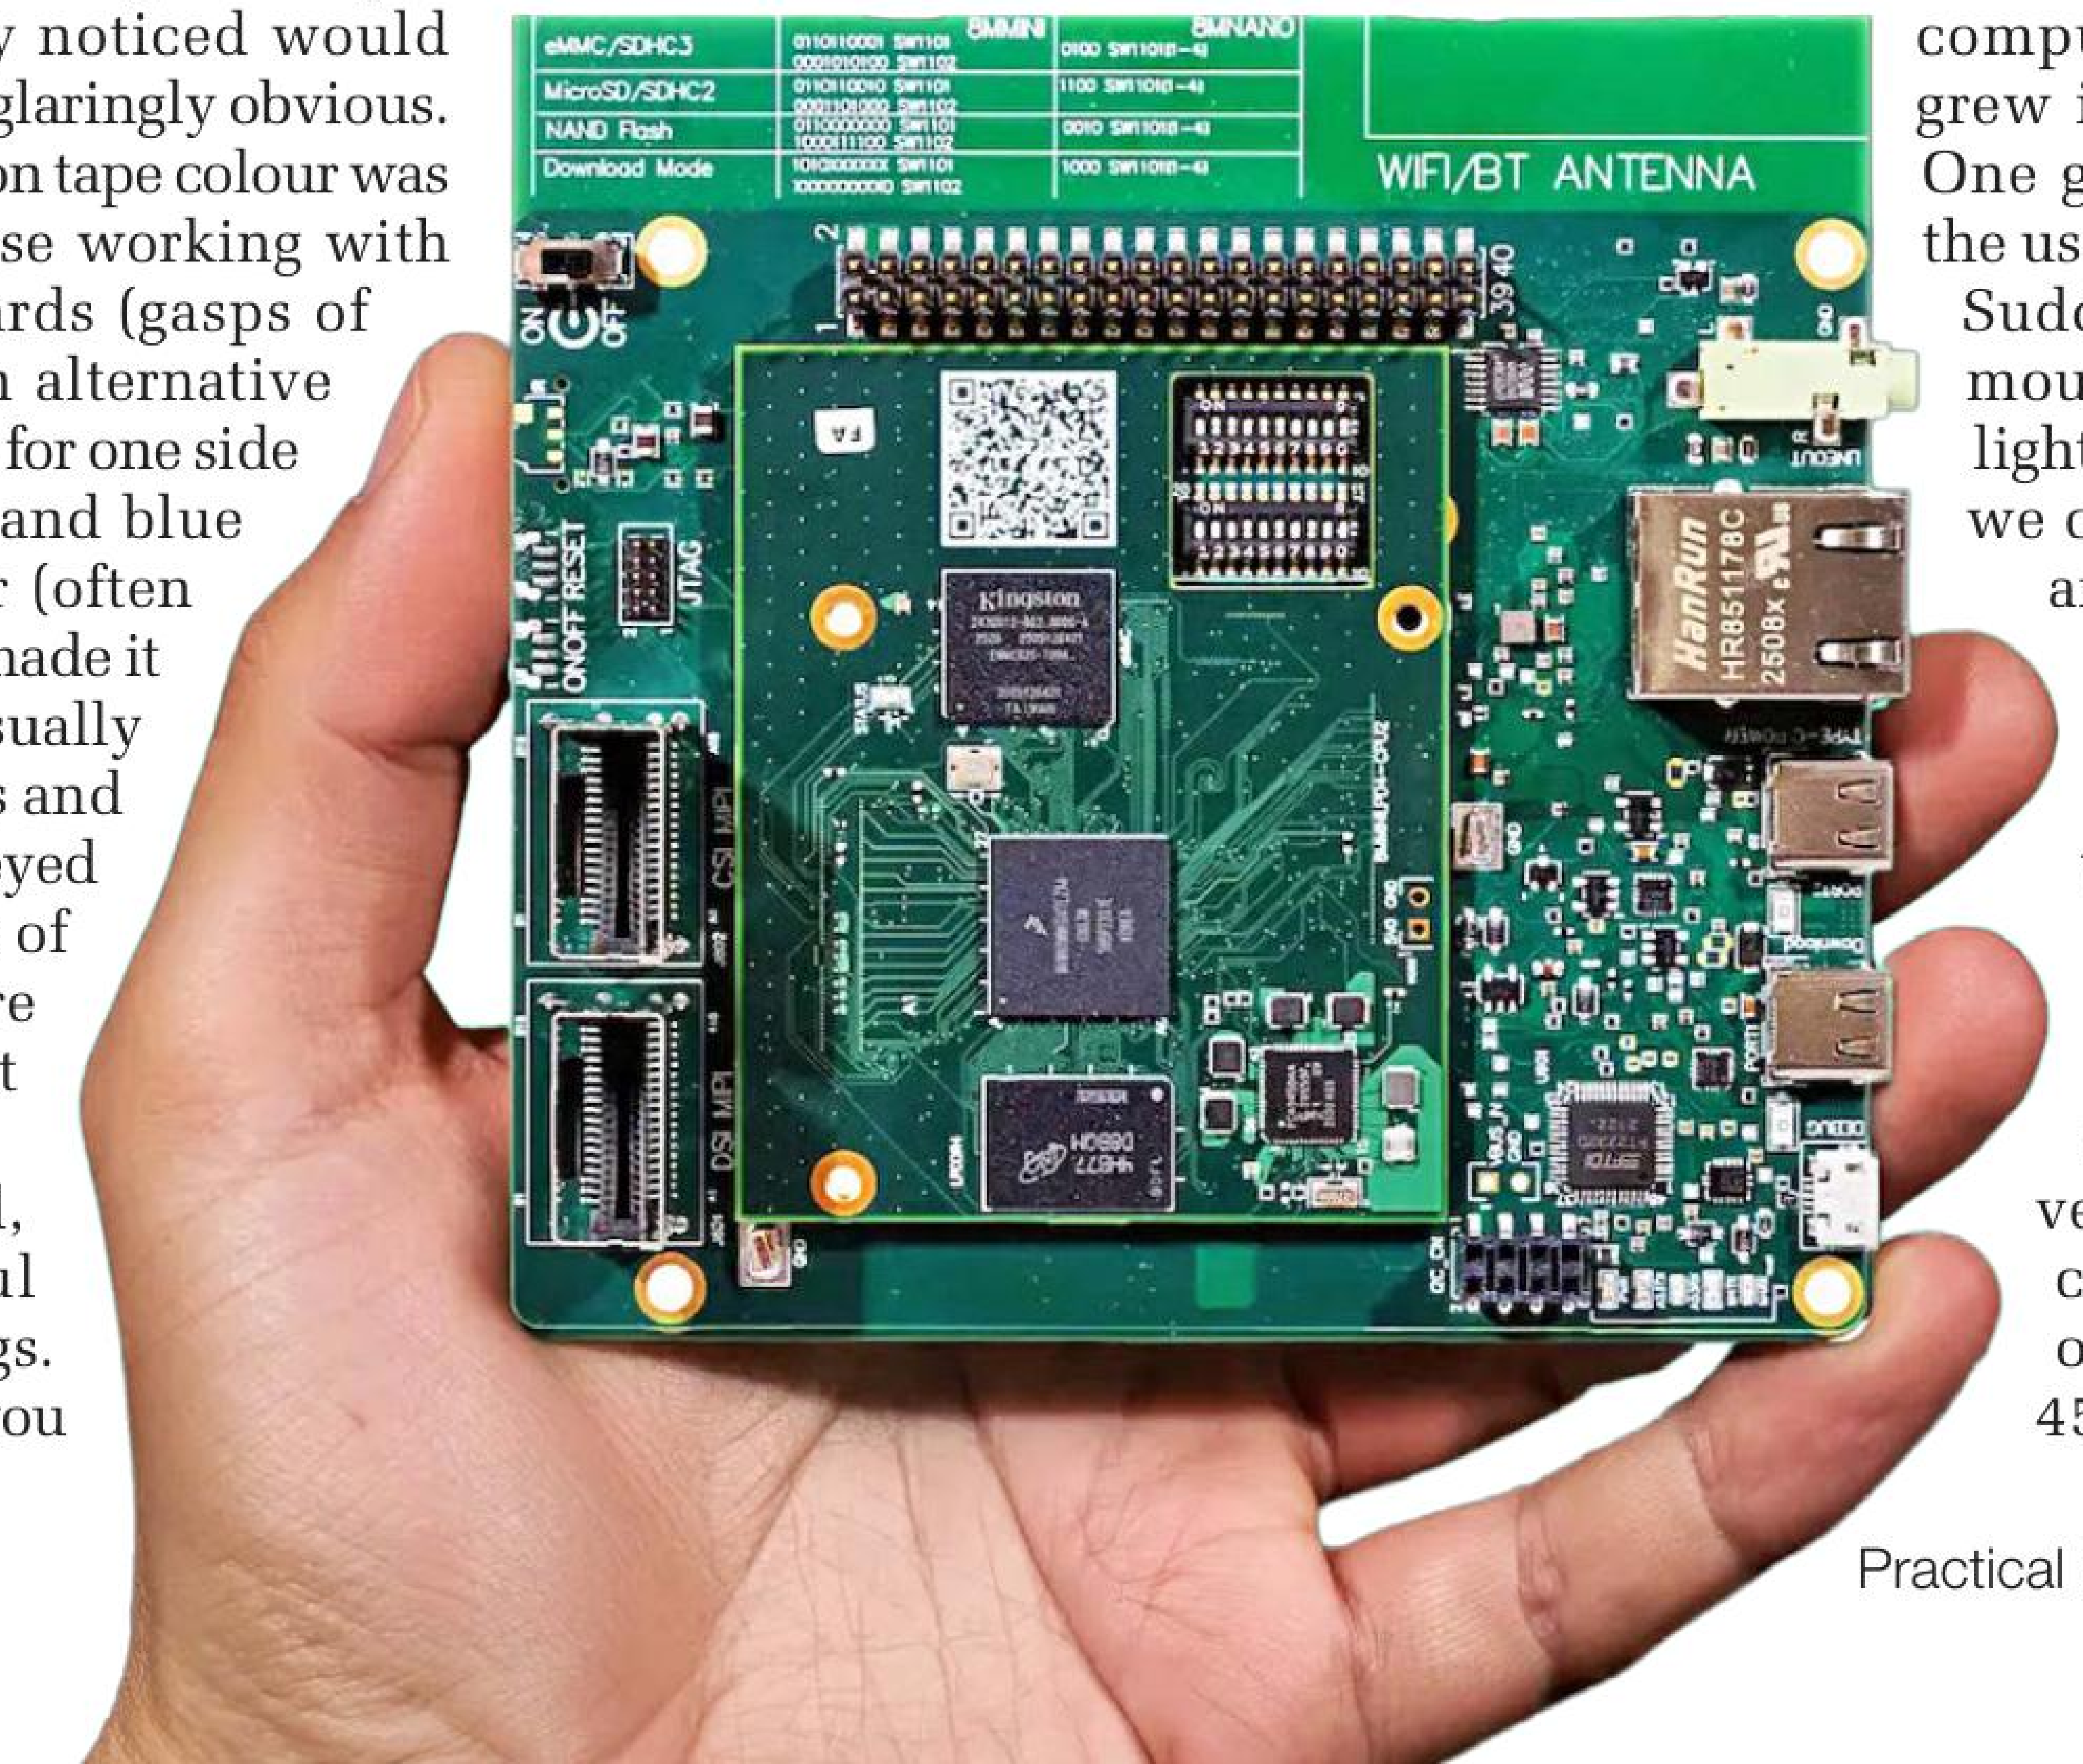

as the tools matured and became more accommodating.

Autorouters began to make an appearance in the late 1970s and early 1980s. By the mid-1980s, they were starting to show up as standard features in PCB layout tools. In theory, these magical constructs would take your netlist and—at the click of a button—produce a fully routed board. In practice… well…

The problem was that early autorouters didn't understand intent, priority, or subtlety; they simply connected point A to point B in whatever way seemed expedient at the time, often leaving behind a trail of unnecessary vias and meandering tracks that looked as if they'd been laid out by an overcaffeinated spider.

As a result, no self-respecting layout designer would admit to using an autorouter, not least because it was almost invariably faster (and certainly more satisfying) to do the job properly by hand than to clean up the enthusiastic chaos that the autorouter left behind.

It was in the latter half of the 1990s that I first saw something that took my breath away. I was chatting with a layout designer who was working on a complex multi-layer board. He demonstrated how he could use his mouse to "push and shove" things like tracks, vias, and components, and how the system automatically re-routed everything around to accommodate these changes on the fly.

I was flabbergasted (my flabber has rarely been so gasted to this day).

Rip-up and retry until it is done

As the 1990s gave way to the early 2000s, autorouting algorithms became increasingly capable. This coincided with signal integrity concerns rearing their ugly heads. Suddenly, it wasn't just about connecting point A to point B; we had to worry about things like controlled impedance, differential pairs and length matching for high-speed signals.

The autorouters rose to the challenge (at least in part) by allowing designers to define constraints and priorities. In practice, however, most experienced layout designers adopted a hybrid approach. The critical nets—clocks, impedance-controlled differential pairs, high-speed buses—were routed by hand, with great care and no small amount of muttering.

Only once the important signals were safely tucked away would we unleash the autorouter on the remaining "mundane" connections, treating it less as a master craftsman and more like an enthusiastic assistant who needed close supervision. The results were certainly better than the overcaffeinated spider efforts of earlier years, but we still kept one hand hovering nervously over the "rip-up and retry" button.

Colour me impressed!

I think we're all aware of the impact of generative AI, starting with the public release of ChatGPT in November 2022. In the past couple of years, several companies have begun offering AI-powered PCB layout tools, with various degrees of success.

Suffice it to say that if layout designers were less than enthusiastic about using autorouters, you can only imagine their sentiments with respect to AI-powered layout.

I have a friend (no, seriously; I do) called Ben Jordan. Ben has decades of experience in PCB layout. He currently works for a company called Quilter (**www.quilter.ai**), which offers a physics-driven, AI-powered PCB layout tool of the same name.

The folks at Quilter are well aware of the scepticism PCB designers have toward AI-powered layout tools. To counter this, they recently did something rather clever. They used Quilter to create a complete, working single-board computer (SBC).

They started by taking an existing open-source reference design from NXP. This design involved a baseboard with a system-on-module (SOM) mounted on top. They defined the outlines and connector placements for both boards, then passed these—along with the schematics and bills of materials (BOMs)—to professional real-world PCB designers, who quoted 200 hours per board (400 hours total).

Next, they fed the same information into Quilter, which completed the SOM in a matter of hours and the baseboard in about a day. Now, this is the clever part, because "the proof of the pudding is in the eating", as they say. They sent the Quilter-generated layouts out for fabrication and assembly.

With the assembled boards back in hand, they first ran tests involving real-world applications with the NXP SOM mounted on the Quilter baseboard. Next, they ran through the same tests using the Quilter system-on-module mounted on the NXP baseboard.

Finally, they replaced some NXP baseboard-SOM combos running in real-world systems with their Quilter counterparts, which passed with flying colours, prompting me to say, "Colour me impressed"!

But wait, there's more!

It's one thing to hear about this sort of thing… it's another thing to see it. I was in the middle of a video conference call with Ben (he was using his main computer, of course) when I mentioned that I'd love to see the Quilter-generated baseboard-SOM combo in action.

Ben responded that he had one on his desk. It already had a camera plugged in because he was testing it with some AI models, so he joined our call a second time via the development board, which left me staring at a brace of Bens. To be honest, I couldn't tell which was which.

Just to add a dollop of metaphorical cream on top of the allegorical cake, while we were still talking, Ben fired up the AI model he was playing with on his Quilter-generated baseboard-SOM combo.

Now, tools like Quilter don't (yet) produce the same level of finely honed, production-optimised layouts that you'd expect from a seasoned human expert with decades of experience.

Where Quilter really shines is in the earlier phases of the design lifecycle. If your goal is to get a working prototype into your hands as quickly as possible—to validate an architecture, test an interface, or simply see if the blasted thing boots—then shaving weeks or months off the schedule is nothing short of transformative.

In this context, "good enough and available tomorrow" can be far more valuable than "perfect and available next quarter".

If history has taught us anything, it's that each new wave of automation has been met with equal parts excitement and scepticism.

AI-powered layout tools like Quilter feel like the next step along that same path. They're not here to replace skilled designers (at least, not any time soon), but they are changing the way we work, especially in the early stages of design. If they can get us from idea to working hardware faster, that's a future I'm more than happy to route toward! **PE**

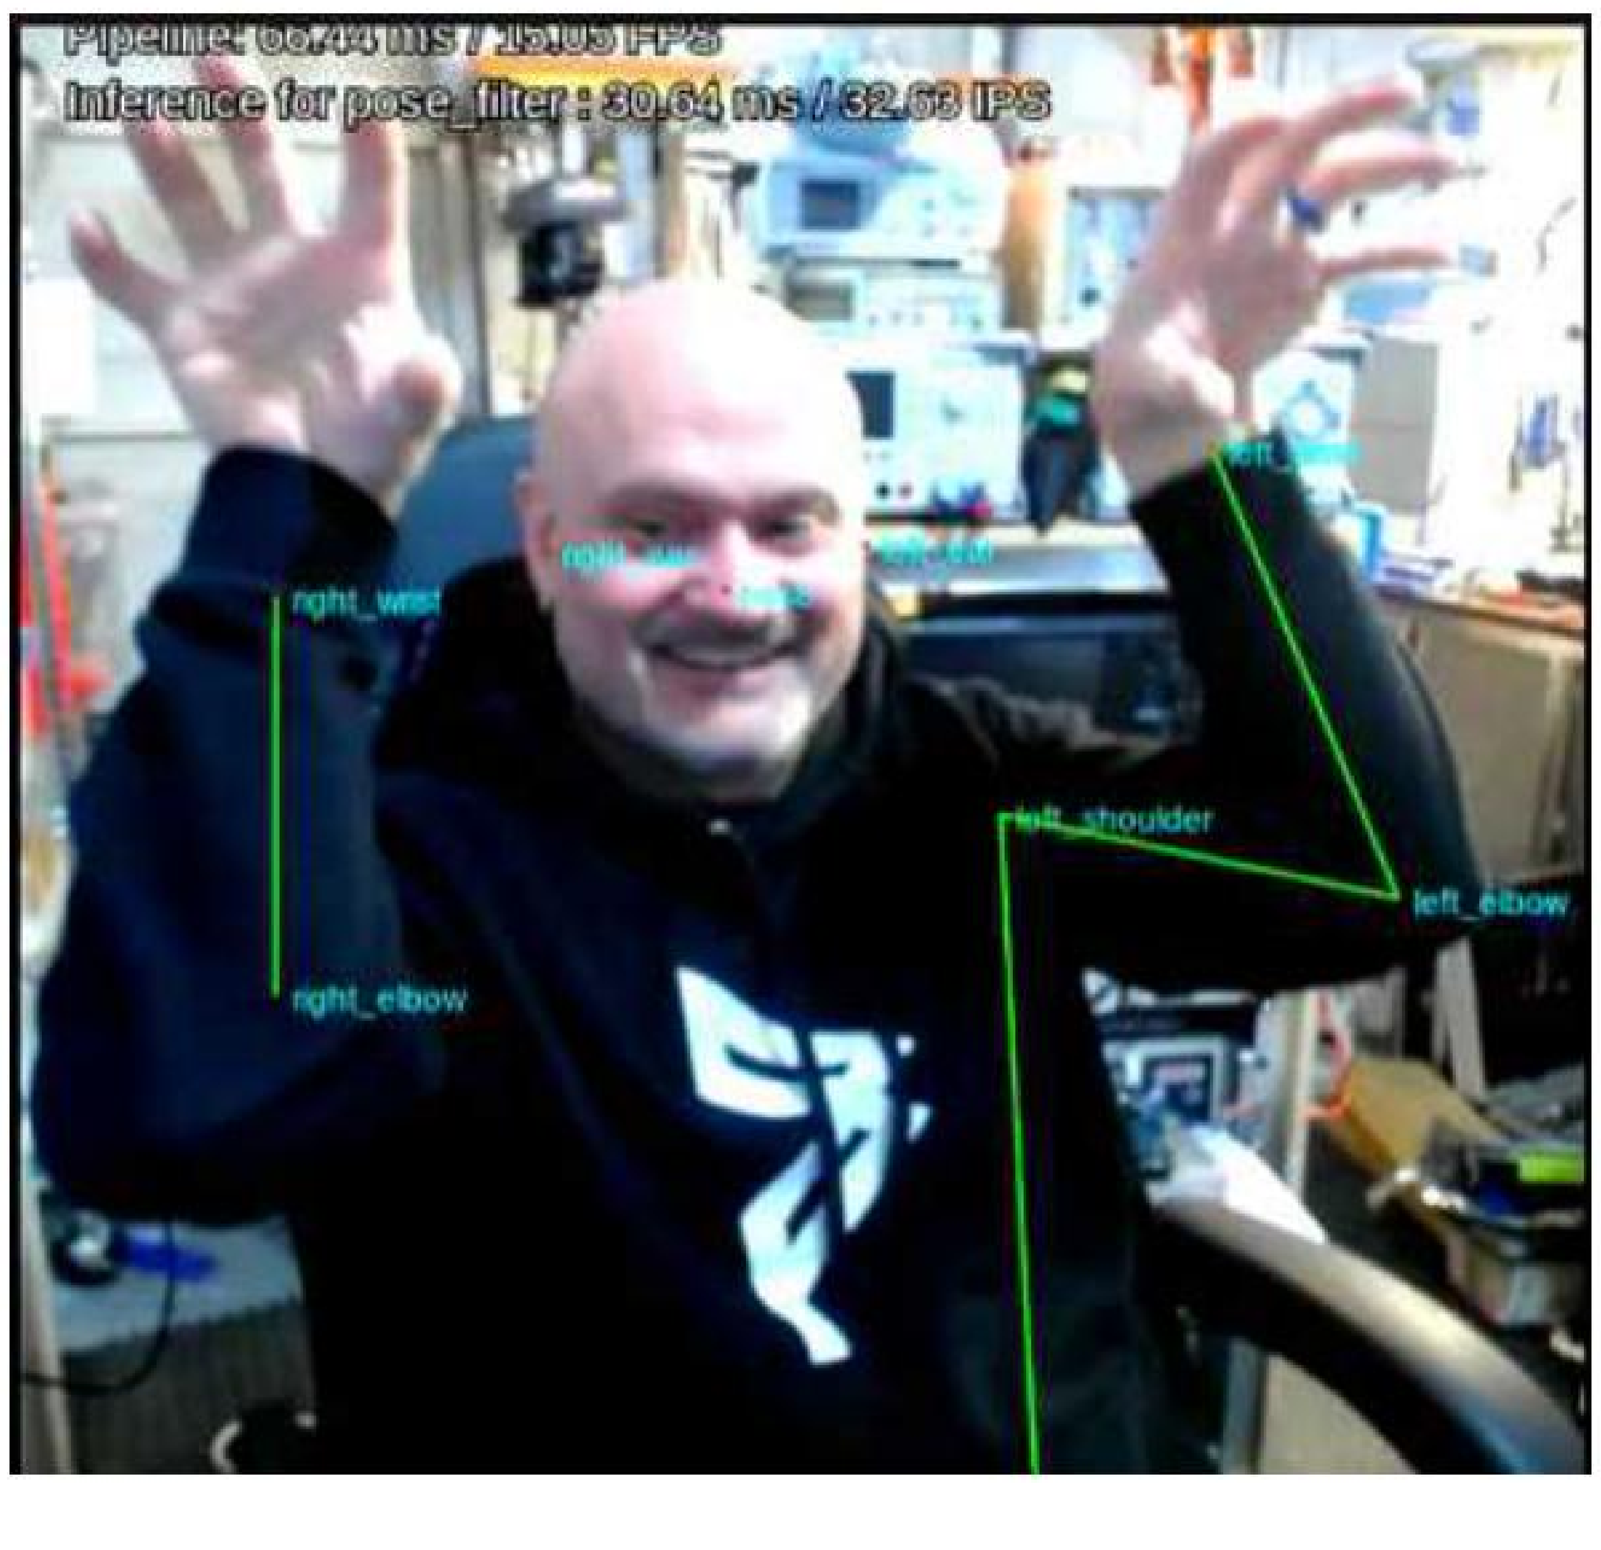

Mains Earthing Systems

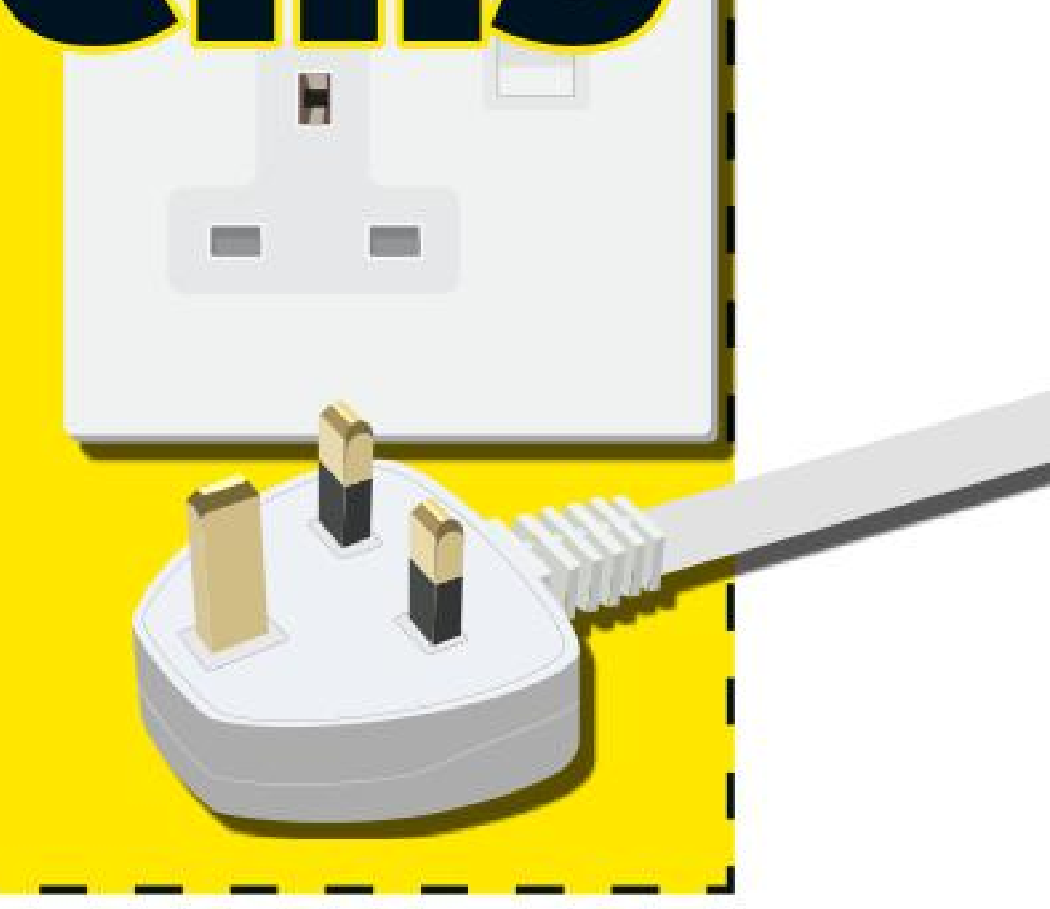

The Earth is an integral part of our power system. It can be used to improve electrical safety, reduce energy losses or save on the cost of a dedicated conductor. Here, we look at the different Earthing systems used worldwide and how they work.

By Brandon Speedie

The Earth's crust is moderately conductive thanks mainly to the salts of sodium dissolved in water and, to a lesser extent, elements such as calcium, potassium, and magnesium. These charge carriers can move freely through soil and rock as long as they remain dissolved in water. The result is a surprisingly conductive electrolyte – see Fig.1.

There are two broad reasons for using the Earth as a conductor.

As a functional conductor

The most obvious use of the Earth is to save on the cost of a dedicated conductor. One example is the Single Wire Earth Return (SWER) line, a common way to distribute power in rural areas. In this case, a significant cost saving can be achieved by only having a single overhead conductor on a power line (see Figs.2 & 3).

The return current (for Neutral) flows through the soil back to the substation or generator. This can sometimes be a distance of hundreds of kilometres/miles.

The Earth is also commonly used in RF applications. A monopole antenna relies on a ground plane to radiate and receive effectively, a role very commonly allocated to 'terra firma' (Latin for "firm land" or perhaps "solid ground").

Another application of Earth is on grid-scale solar farms. Solar panels are effectively three-terminal devices; they have a positive output, a negative output and a frame or chassis.

If the frame is left electrically unconnected, it can float to a different voltage from its other two terminals. Charge carriers will then begin to migrate out of the substrate in a process known as Potential Induced Degradation (PID). This leads to reduced yield and eventually, early failure.

On commercial solar farms, care is taken to ensure the panel mounting solution is well bonded to Earth and that the array DC voltage does not float too far from the Earth's potential.

Safety

Perhaps the most prominent function of Earth is to provide electrical safety.

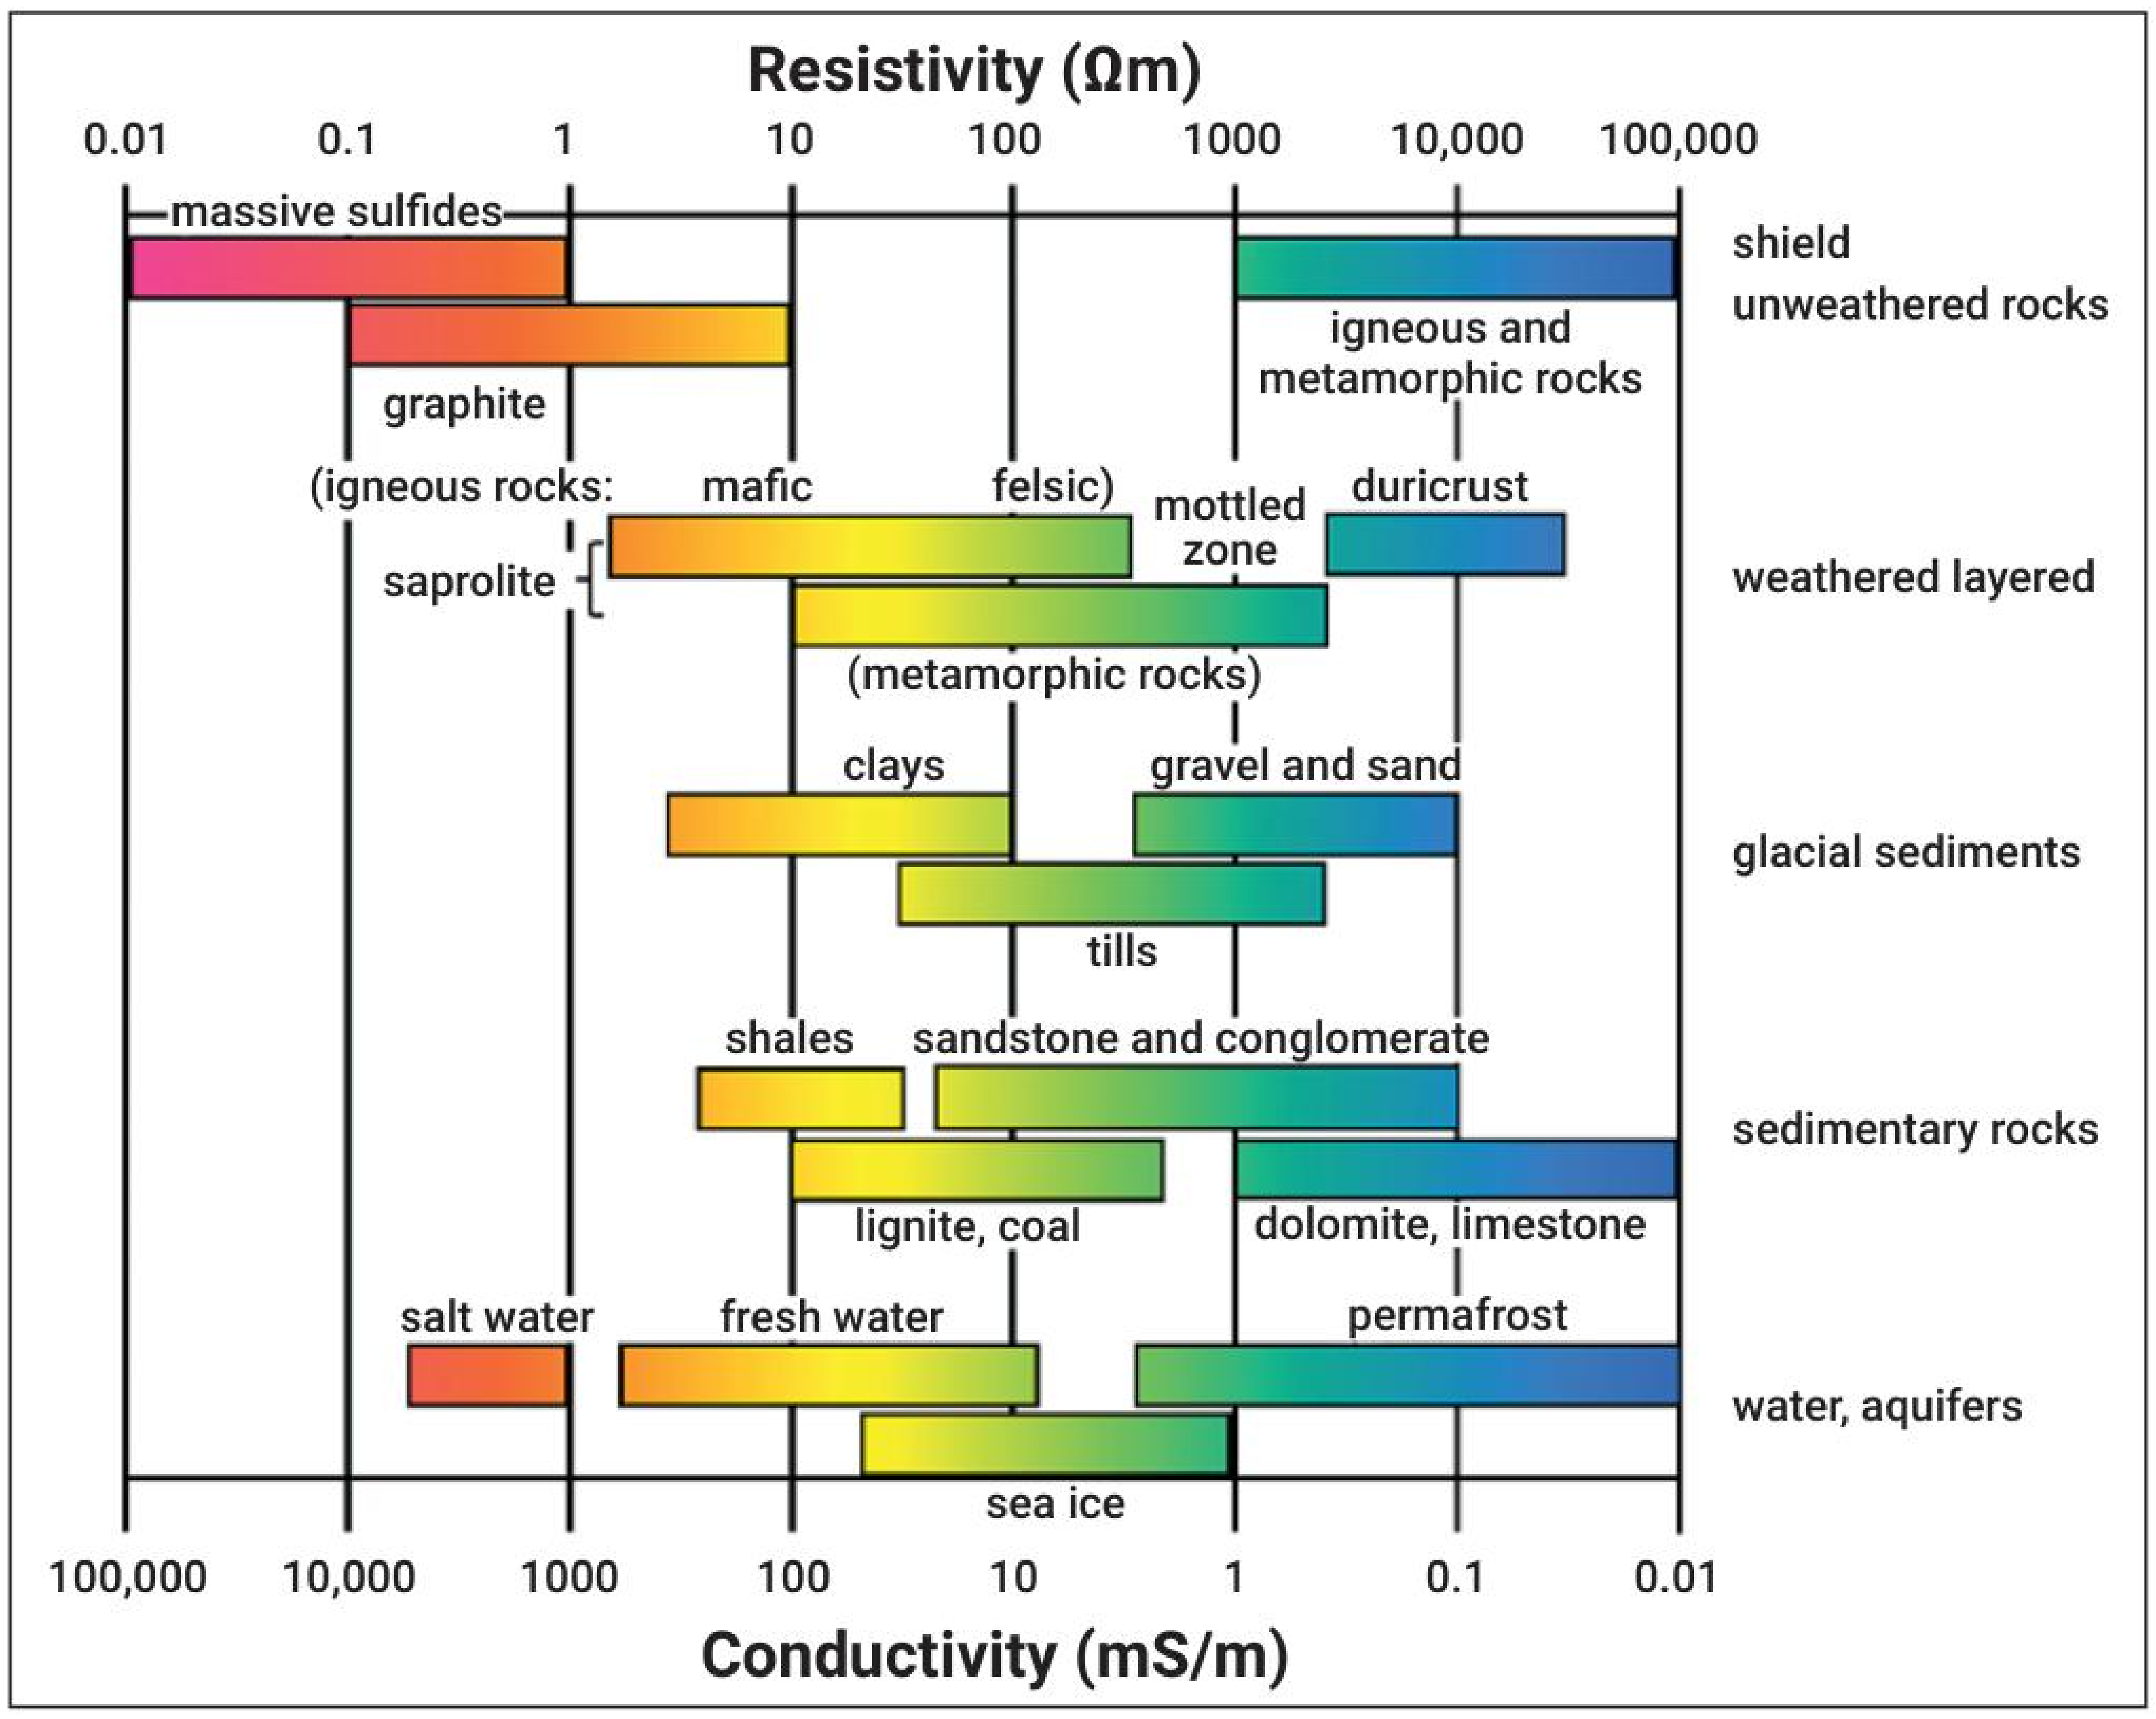

Fig.1: resistivity figures for some common components of the Earth's crust. Note the different units on the top and bottom horizontal axes, which are inversely equivalent; as S (siemens) is the inverse of Ω (ohms), 1mS is equivalent to 1kΩ. ***Source: GeoSci Developers – https://pemag.au/link/abu7 (CCA 4.0).***

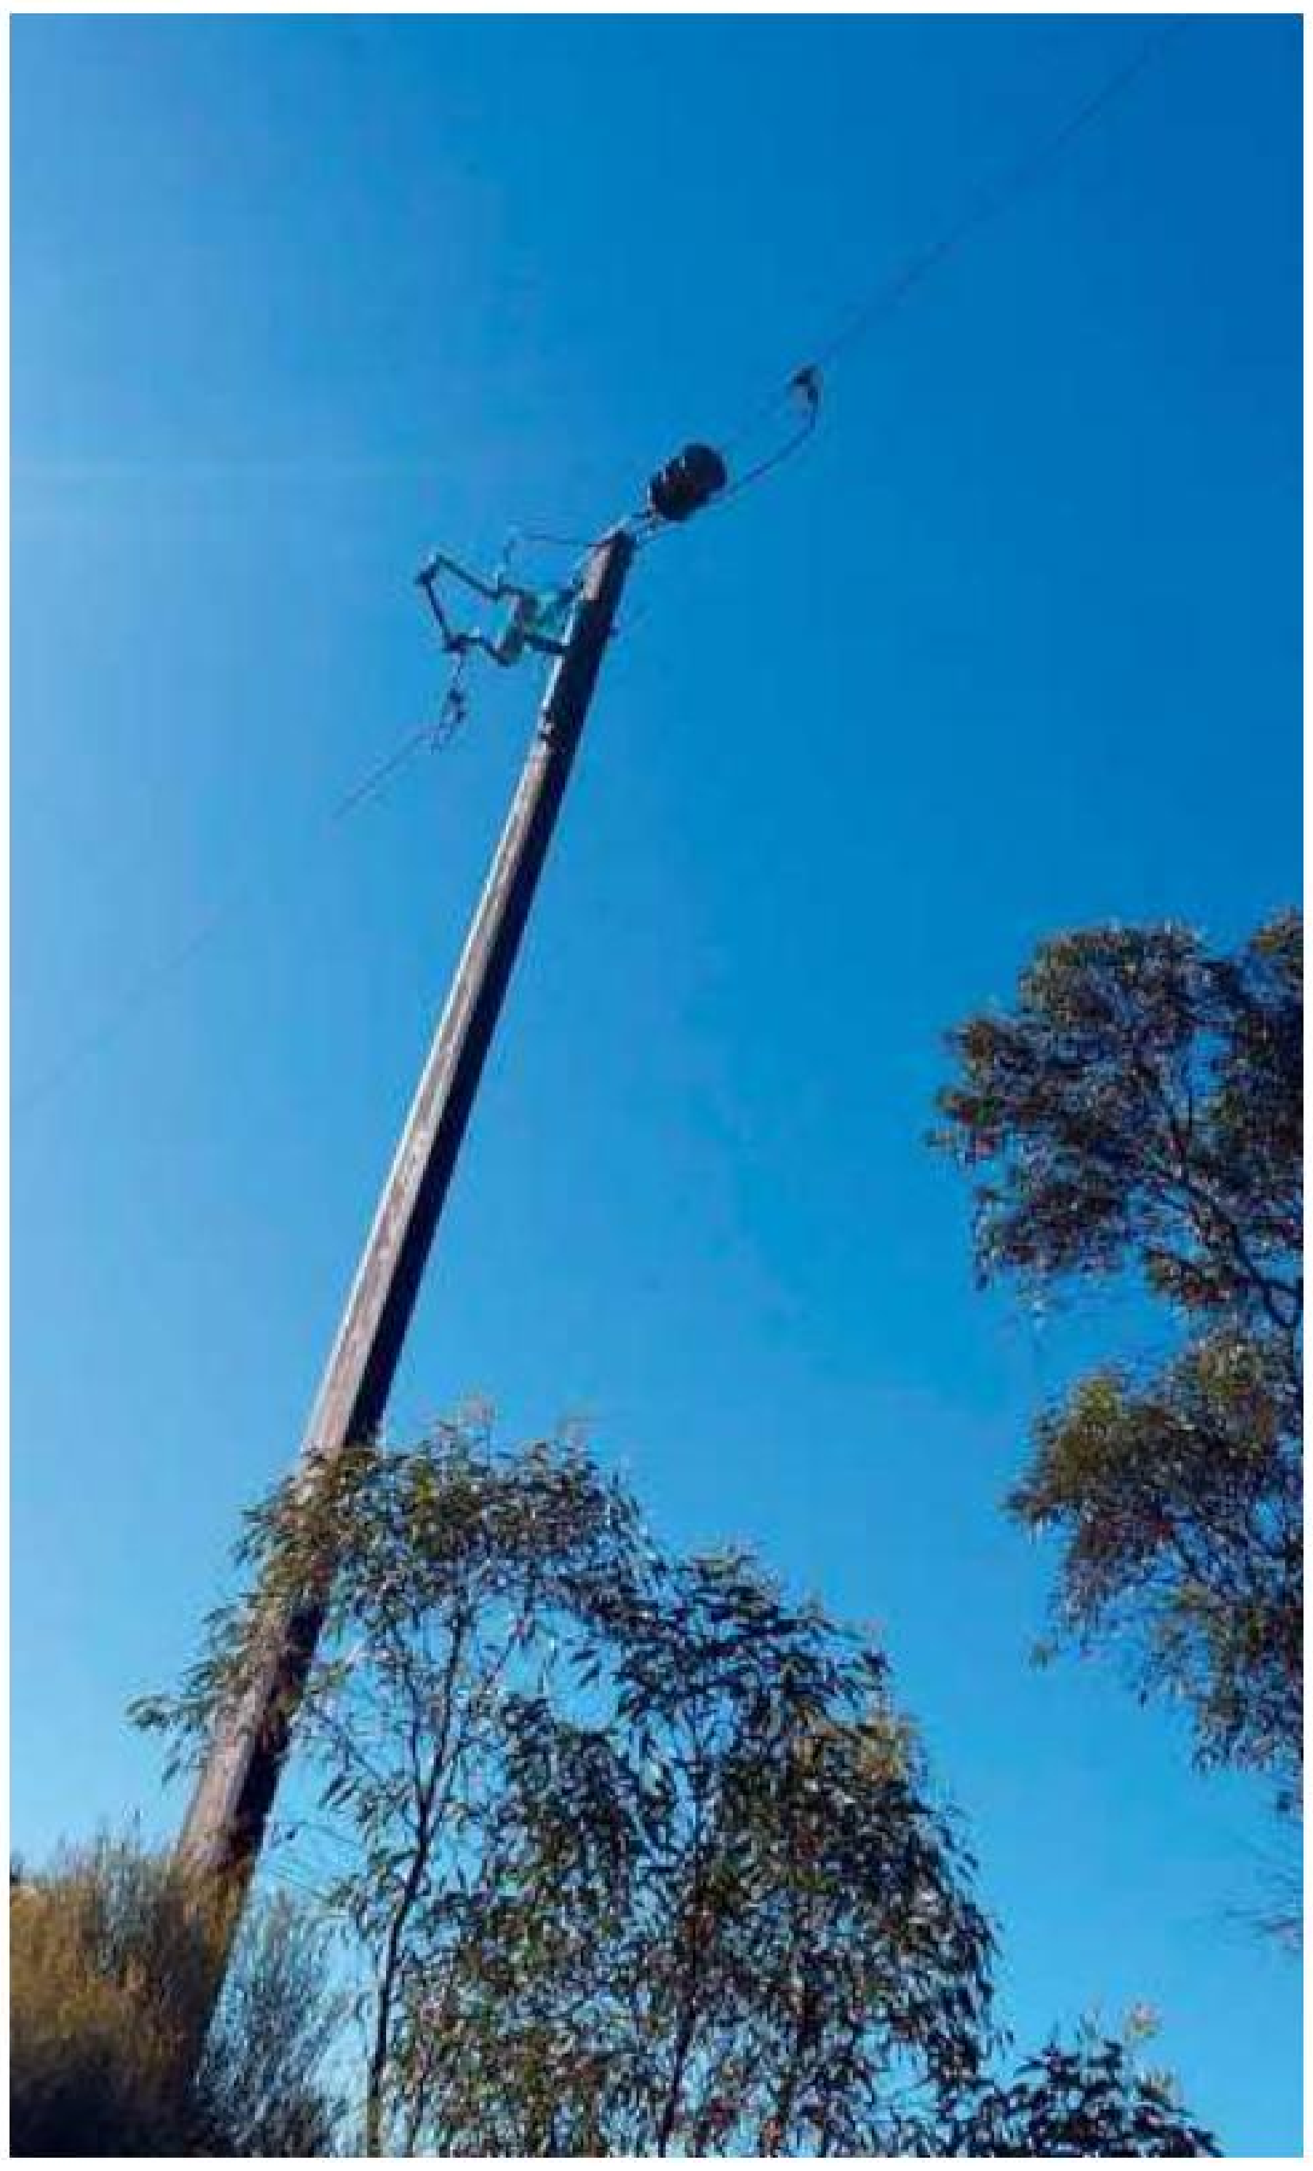

Fig.2: SWER line in South Australia. The unusual pole construction is concrete sandwiched between two steel beams, known as a "Stobie pole".

Mains Earthing Systems

In normal operation of a single-phase AC circuit, current flows into or out of the Line conductor, through the load, and returns via the Neutral conductor. In a fault scenario, an Earth connection gives a low impedance path for current to flow, which will usually trip a circuit breaker.

In some scenarios, the fault will not draw enough current to trip the circuit breaker, but it should be enough to trigger a Residual Current Device (RCD).

In an RCD, the Line and Neutral conductors both pass through a current transformer (CT). In the absence of a fault, the current flow is balanced between Line and Neutral, so the magnetic fields of these two currents cancel, and no net current is detected – see Fig.4.

In a fault scenario, current flows through Line, but not all is returned via the Neutral; some flows through the Earth connection.

This imbalance is detected in the RCD, which will typically trip once the imbalance exceeds 30mA (although more sensitive RCDs exist, eg, 15mA; the trip current is a balance between sensitivity and nuisance tripping).

The Earth can also be used to ensure electrical safety during the normal operation of a grid. The most prominent such application is lightning suppression. If the potential difference between the Earth's surface and the power lines were left uncontrolled, a direct strike from a lightning bolt would charge up the network to a high voltage, leading to arc-over at the insulators.

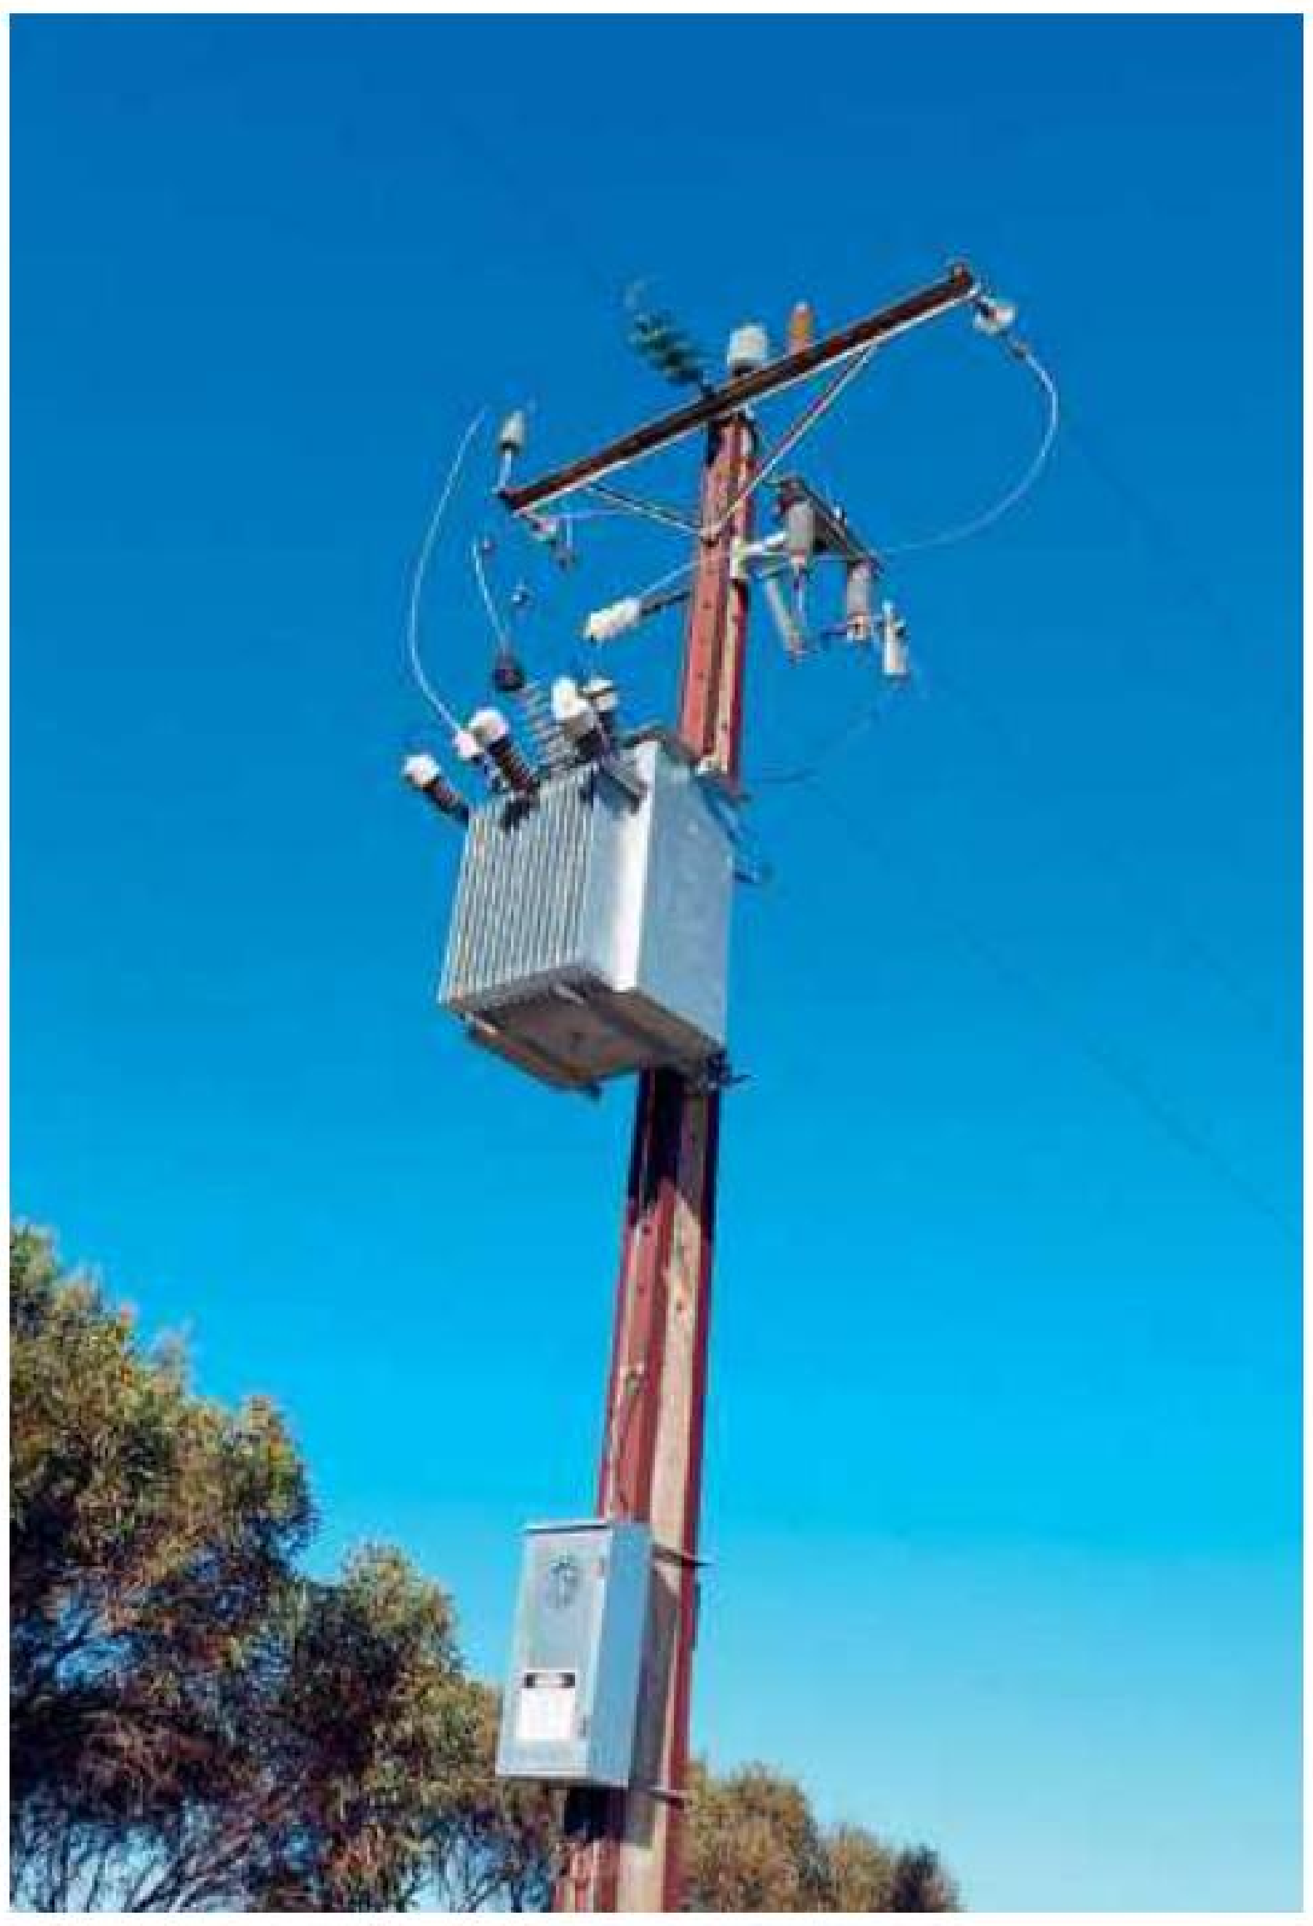

Fig.3: the start of the SWER line shown in Fig.2. The three conductors on the right are 33kV phase-to-phase, or 19kV from phase to Neutral/Earth. The SWER line taps off the middle phase and extends to the left. The return Neutral current flows via Earth.

It is therefore critical that this energy is shunted to Earth to maintain grid tolerances. Earth is also a logical place to shunt this charge as the lightning originates from a static buildup between the ground and the atmosphere.

Types of Earthing systems

Earthing schemes used in a mains grid are commonly described by a sequence of letters based on where the circuit Earth originates from.

"T" (terra; Latin for "Earth") refers to a direct connection to the soil of the Earth. This is typically achieved by driving a conductive stake into the ground, or perhaps multiple stakes and/or bonding to buried metal pipes. In larger installations, such as substations or generators, a dedicated buried circuit or 'Earthing ring' made of bare wire (usually copper) encircles the installation.

"I" (insulatum; Latin for "insulated") means no connection to Earth or a high-impedance connection through an Earthing resistor.

"N" (network) means the Earth connection is via the network or grid. Network Earths will still connect to the soil at some point, but this may be some distance away, not at a local Earth stake, as with terra (T).

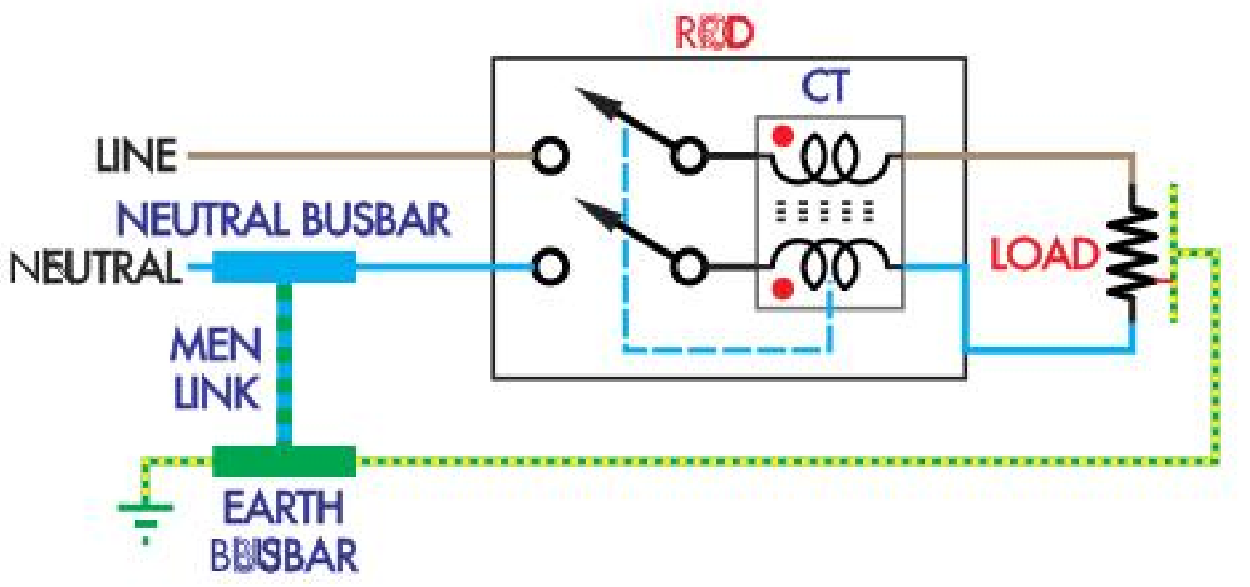

Fig.4: a Residual Current Device (RCD). Usually, current flows to or from the Line conductor through the load and is returned via Neutral. The magnetic fields of the two conductors are cancelled, so the CT detects no net current. A small amount of leakage between Line and Earth, shown as a thin red line here, is enough to trip the RCD.

"C" (combined) means the circuit Neutral and Earth are combined into a single conductor in the network.

"S" (separate) is where the circuit Neutral and Earth are run as separate conductors in the network.

The Earthing system can thus be described by two letters, the first indicating the source Earth, and the other the load Earth.

TT (Terra-Terra)

Terra-terra networks are physically connected to Earth at both the generator and load (see Fig.5). Typically, this will be at the low-voltage distribution transformer and the customer's premises.

TT networks rely heavily on the Earth connection's integrity, so care is taken to ensure a sufficiently low Earth loop impedance. This can include tight specifications around Earth stake

Terminology

Phrases such as Earth, Neutral, common and ground are sometimes used somewhat interchangeably. They can be confusing terms, particularly from our perspective as electronics enthusiasts.

Earth: a connection to terra firma, either directly through an Earthing stake, or via a conductor that is bonded at some point with Earth.

Neutral: the return current of a single-phase AC supply. Typically, this will be the centre point of a star-connected three-phase circuit. In regular operation, most networks should have very little voltage difference between the Neutral and Earth.

Ground: a common node in a circuit, usually at 0V DC potential. Confusingly, ground can be 'grounded' by tying to 'Earth', but it is uncommon in modern usage. Floating circuits are generally considered to have a 'ground', but it could drift relative to Earth; it is usually the negative end of a battery or similar and is used as a local reference and/or current return. In a circuit running directly from the mains, 'ground' may even be connected to (or very close to) the Line potential! Some circuits can have multiple grounds (analog, digital etc).

Common: a node in a circuit shared by many components. It is sometimes used interchangeably with 'ground' but can also be used where multiple signals are tied together. Examples are a common bus in a multiplexed display or a common signal tying multiple opto-isolators together.

construction and placement, as well as considerations of soil conductivity.

This is particularly important in cold areas, where frozen soil dramatically increases resistivity, or in high rainfall regions, where soil electrolytes are diluted, leaving few charge carriers for conduction.

Even given these additional requirements, a standard overcurrent breaker is not guaranteed to trip from an Earth fault. As a result, TT customers will almost always need Earth leakage protection in the form of an RCD.

Historically, TT networks were not popular due to the difficulty of ensuring safety without Earth leakage protection. The advent of cheap RCD breakers has led to its increasing use worldwide, such that it is now the most common scheme.

Many parts of Europe, including France, Denmark, Belgium, Spain, Italy, and Portugal, are now predominantly TT, as well as Japan, Malaysia, Argentina and many others. Germany extensively uses TT outside of metropolitan areas.

IT (Insula-Terra)

Insula-terra networks are connected to Earth at the customer's end but not at the generator or distribution transformer (Fig.6). Therefore, the Line and Neutral connections have no reference to Earth, which minimises shock hazards. The reader may recognise this advantage from using an isolation transformer when working on mains-powered electronics.

IT networks are often referred to as "first fault free", as any fault will convert the system into another scheme (usually TN) while the fault is present, and subsequent faults may be dangerous. This is why IT networks are not common worldwide, except in specialised applications.

This includes hospitals, where patients are at a higher risk of shock when coupled to medical equipment, and industrial areas where a flammable atmosphere may be present, so the risk of sparking needs to be minimised.

Scandinavia is an exception, where frozen ground and rocky geology make Earthing difficult. Norway in particular makes heavy use of IT Earthing.

In India, a variation of the IT network called Resistance Earthed Neutral (REN) is used in mining areas. A Neutral grounding resistor limits the Earth current to 750mA.

TN-C (Terra-Network-Combined)

Terra-network-combined systems get their Earth from the network by combining it with the Neutral conductor (Fig.7). This combined conductor is commonly referred to as the Protective Earth Neutral (PEN).

In a TN-C network, the distribution transformer is Earthed at the source end, which is also connected to the circuit Neutral. This PEN conductor then runs along the poles and wires of the grid to customer premises, where it is used as the "Earth".

TN-C networks do not require an Earth stake at the customer premises or RCD breakers as in a TT or IT network, but they are extremely reliant on the integrity of the PEN. If there is a break in this conductor, the customer load will act like a pullup resistor, raising the potential of Neutral/Earth to mains voltage; a hazardous situation (see Fig.11).

TN-C networks also suffer from conducted interference. As the circuit Neutral is combined with the Earth, coupled noise from heavy industrial equipment can pass through the network and cause problems with sensitive equipment such as telecommunications broadcast infrastructure.

TN-S (Terra-Network-Separate)

Terra-network-separate networks run a dedicated Earth on the network, separate from the Neutral conductor (Fig.8). The distribution transformer is Earthed at the source end and connected to two conductors. One is the Neutral, while the other is a dedicated Earth – see Fig.13. TN-S networks are the safest configuration but are also more expensive, given the added conductor. TN-S is used extensively in India.

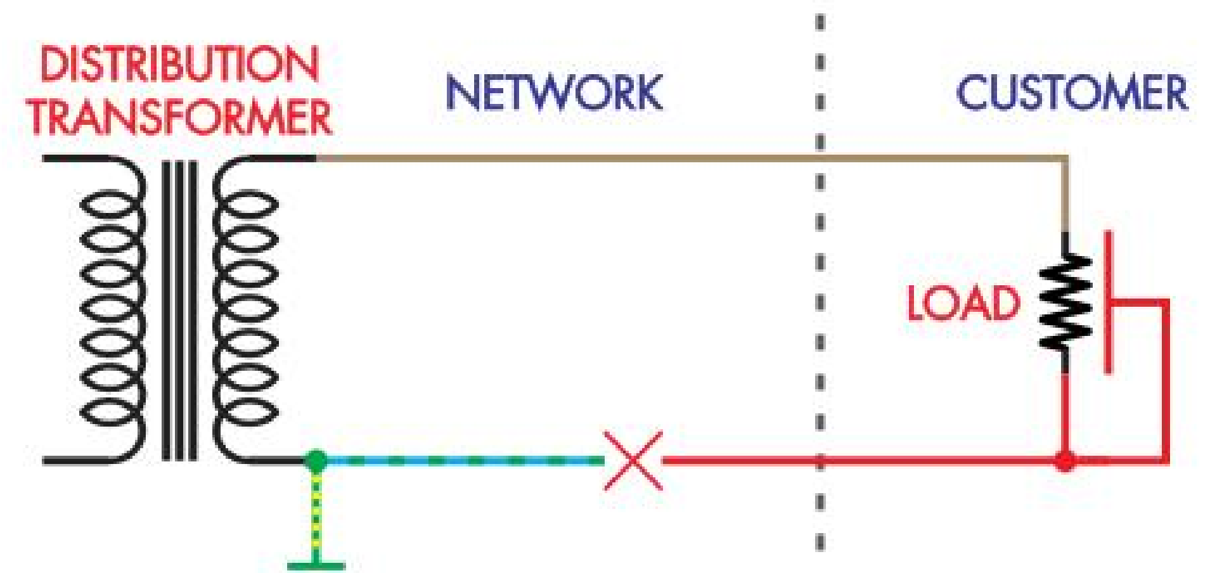

Fig.11: a Neutral fault with the TN-C scheme. The customer load acts like a pullup resistor, raising the Earth to a high voltage and creating a shock hazard.

TN-C-S (Terra-Network-Combined-Separate)

Terra-network-combined-separate networks are a hybrid of the TN-C and TN-S systems. The source transformer is Earthed, while a combined PEN conductor emanates onto the network (Fig.9). The PEN is split into dedicated Neutral and Earth conductors at some agreed location (usually the customer's switchboard).

TN-C-S is widely used in the UK, USA, Canada, Israel, Australia and New Zealand – see Fig.14. Germany also predominantly uses TN-C-S in metropolitan areas.

MEN (Multiple-Earthed-Neutral)

Australia and New Zealand use the Multiple Earthed Neutral (MEN system) – see Fig.10. It is a TN-C-S system, though TT may be permitted in some situations – usually rural areas. Unusually for TN-C-S, an Earth stake is mandatory at each customer premises.

In MEN networks, the source distribution transformer is Earthed, and a combined PEN conductor runs in the grid. The PEN is Earthed at multiple points throughout the network, including at the customer stake.

This gives good immunity against a broken Neutral; if the customer is downstream of the fault, their PEN will not rise to a dangerous potential thanks to the Earth stake at their premises and any neighbouring premises or network Earths.

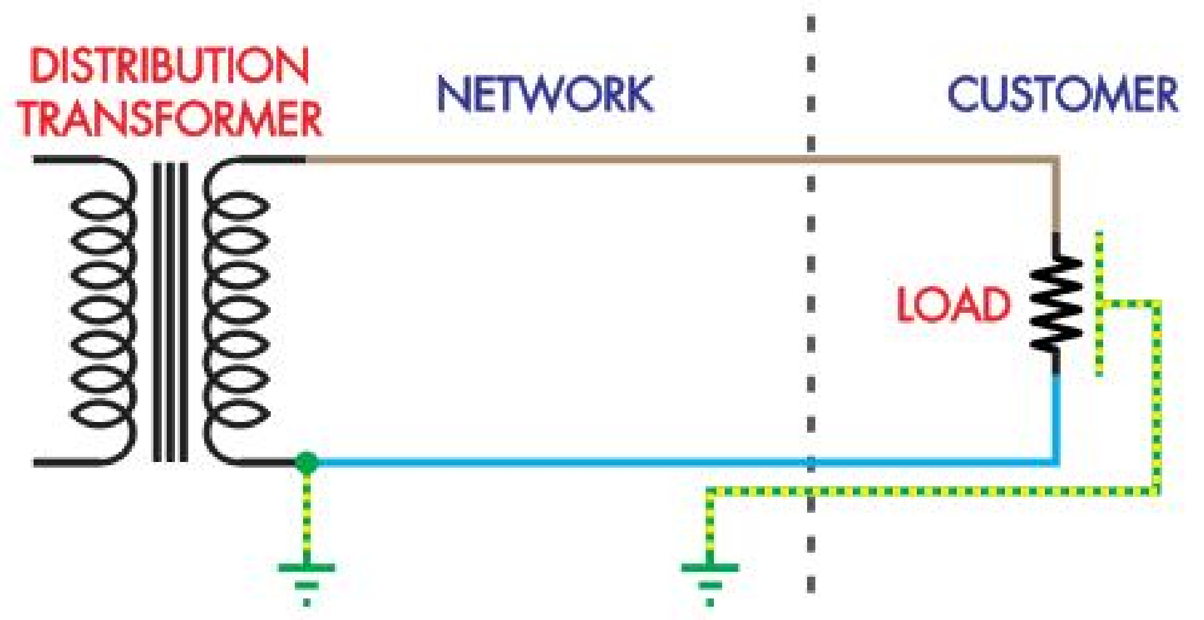

Fig.5: the TT Earthing scheme.

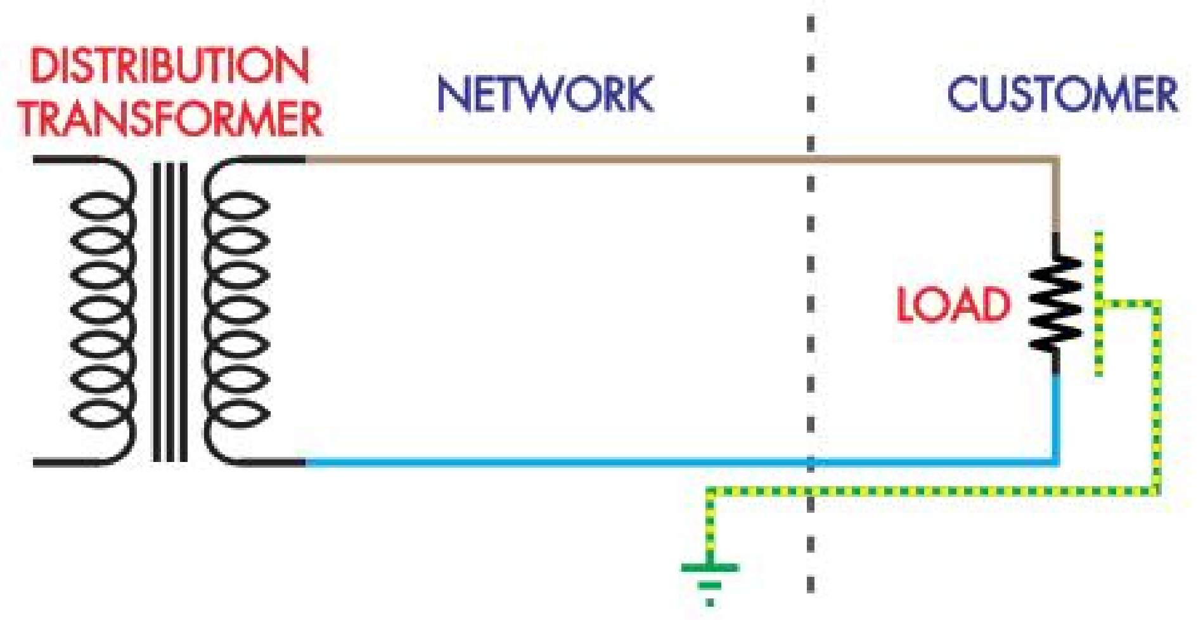

Fig.6: the IT Earthing scheme.

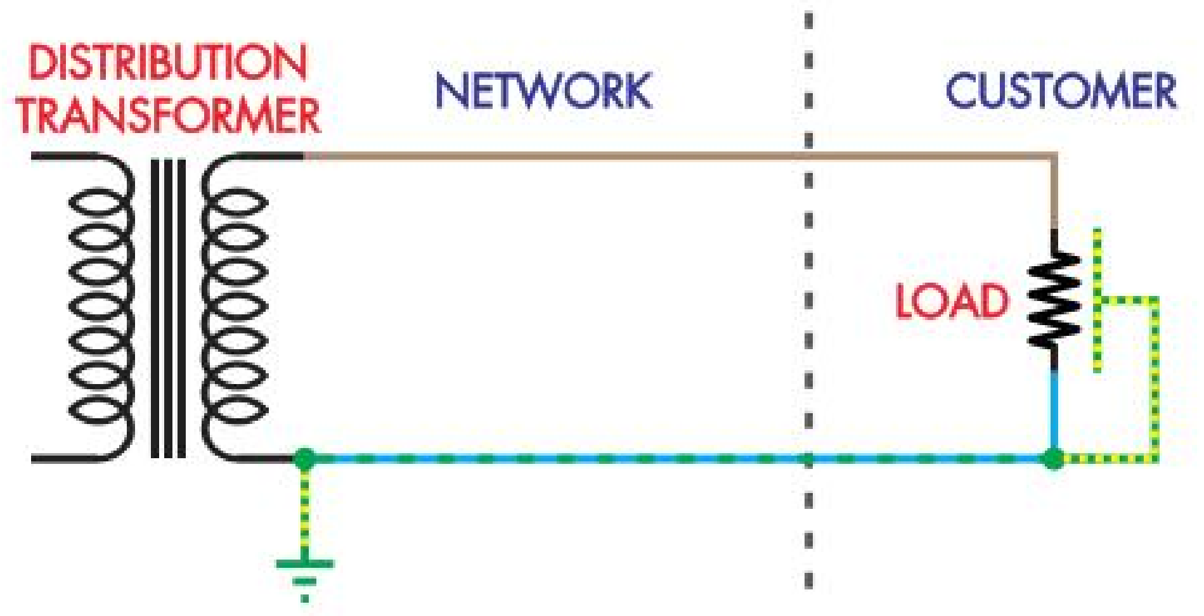

Fig.7: the TN-C Earthing Scheme.

Mains Earthing Systems

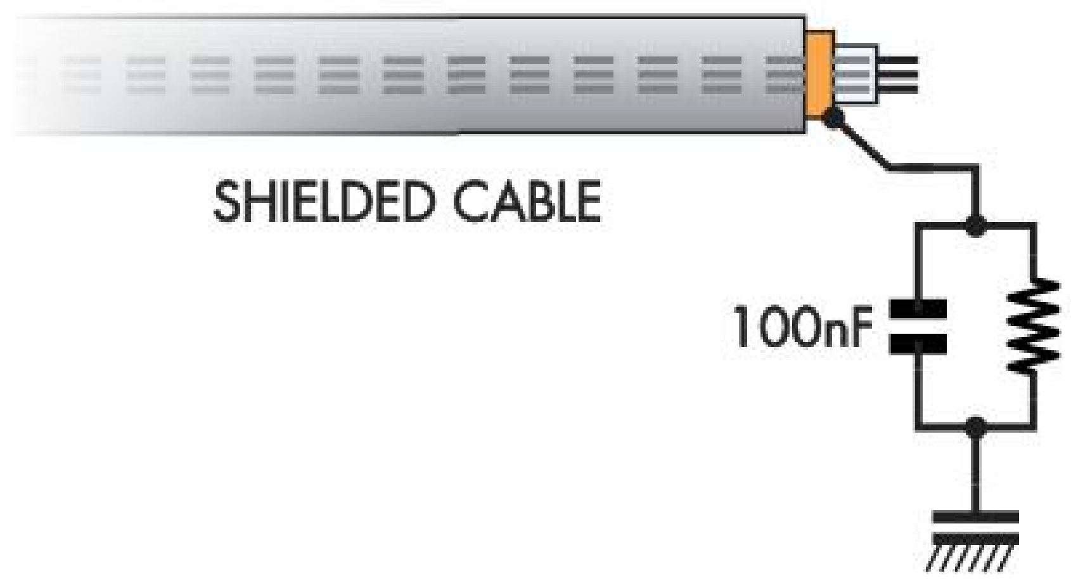

Fig.12: this configuration stops the communications cable shield from drifting too far from Earth. It keeps it at a low AC impedance via the capacitor but will not form an 'Earth loop'.

At the customer premises, the combined PEN connects to the Neutral bar in their switchboard. This busbar then distributes the Neutrals to all of the circuits within the installation. Separately, a dedicated Earthing busbar connects to the "Earth" conductor emanating throughout the property, as well as the Earth stake.

The Earth busbar and the Neutral busbar are joined by a single strap, known as the "MEN link". This link is the separation point between the TN-C scheme on the network and the TN-S scheme within the customer premises.

Earth integrity

The Earth connection is convenient, as it can be assumed to be the same voltage across multiple installations, even if they are geographically separated. But as the soil has a finite resistance, this is not always true.

This can create problems where two circuits are bonded to mains Earth and are also linked by conductors, similar to a local 'ground loop' or 'earth loop' that can cause problems with audio electronics.

This is particularly problematic when large conductive structures are located near high-power switching gear. This might be a steel fence around a switchyard or a buried gas pipeline adjacent to power lines. In a fault, a large current might flow into 'Earth', raising the local voltage near the Earthing stake/ring. If the metal structure is close enough to this fault, its voltage will rise.

If someone touches this structure, they may receive a lethal voltage despite being far from the actual fault. This is because the metal is a better conductor than the Earth, so the ground they are standing on is at a different voltage than the ground near the fault. This is known as a 'touch potential' and is a major hazard in high-power assets.

Editor's Note: in extreme cases, there can be enough potential between workers' feet to electrocute them. Electrical workers are trained to hop if they suspect such a fault exists!

Similarly, industrial Ethernet networks can also suffer from ground loops and unequal Earthing. Ethernet uses differential signalling, so it is commonly run over UTP (Unshielded Twisted Pair) cabling. A high Common Mode Rejection Ratio (CMRR) amplifier cancels any coupled noise or interference at the receiver. Thus, shielded cable is not needed for noise immunity.

Despite this, shielded or foiled Cat6 cabling is common in industrial settings. Often, the designer will reason that shielded cable will be better than unshielded, so it is worth the marginal cost increase. However, it can often be more trouble than it is worth.

Ethernet uses "magnetics" (signal transformers) at the receiver and transmitter to galvanically isolate the channel, preventing ground loops from forming. Shielded cable effectively breaks this isolation by connecting a conductor directly between the receiver and transmitter.

If there is any Earth imbalance between this equipment, large currents can flow, which can cause interference or damage. For this reason, if shielded Ethernet cable is used, it is often only Earthed at one end of the cable, or better still, connected through a parallel RC combination, perhaps 1MΩ || 100nF (see Fig.12).

The resistor weakly holds the shield at a known voltage, while the capacitor offers a low impedance path for AC voltages (coupled interference).

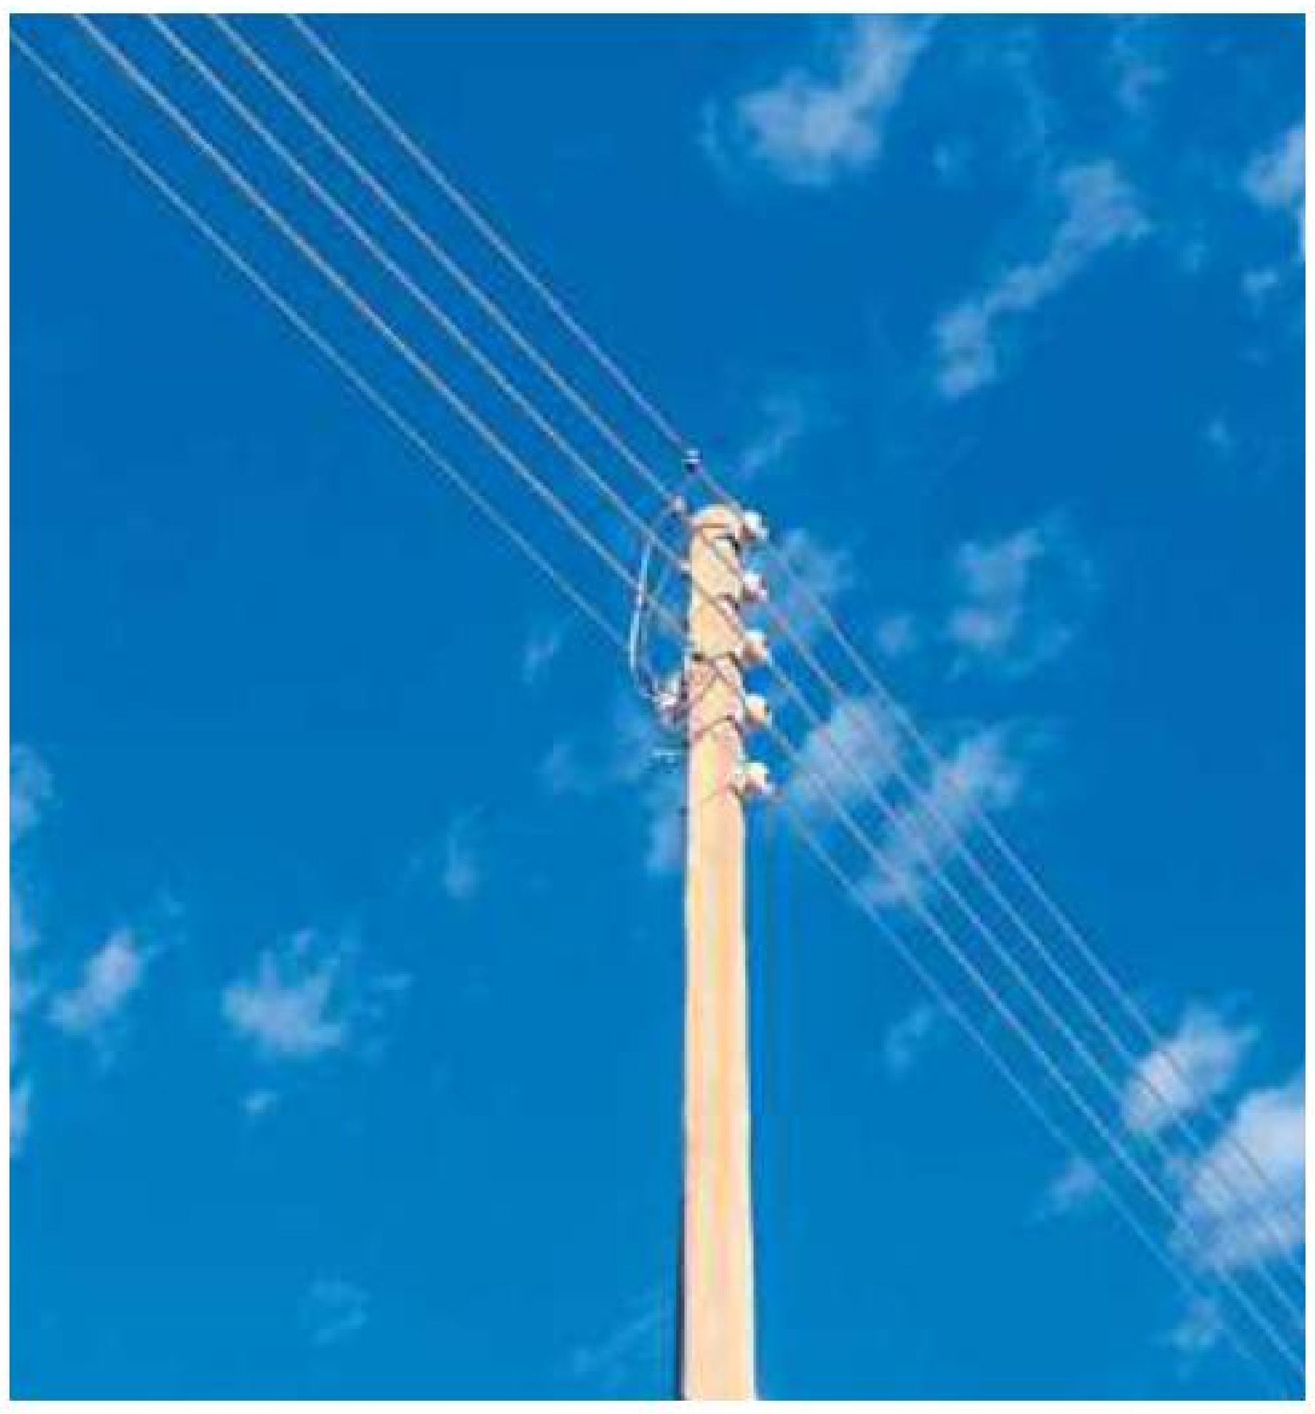

Fig.13: the TN-S scheme in Namibia. The five conductors (from bottom to top) are Earth, Neutral & three Line phases. Note how the Earth wire has a smaller diameter than the others.

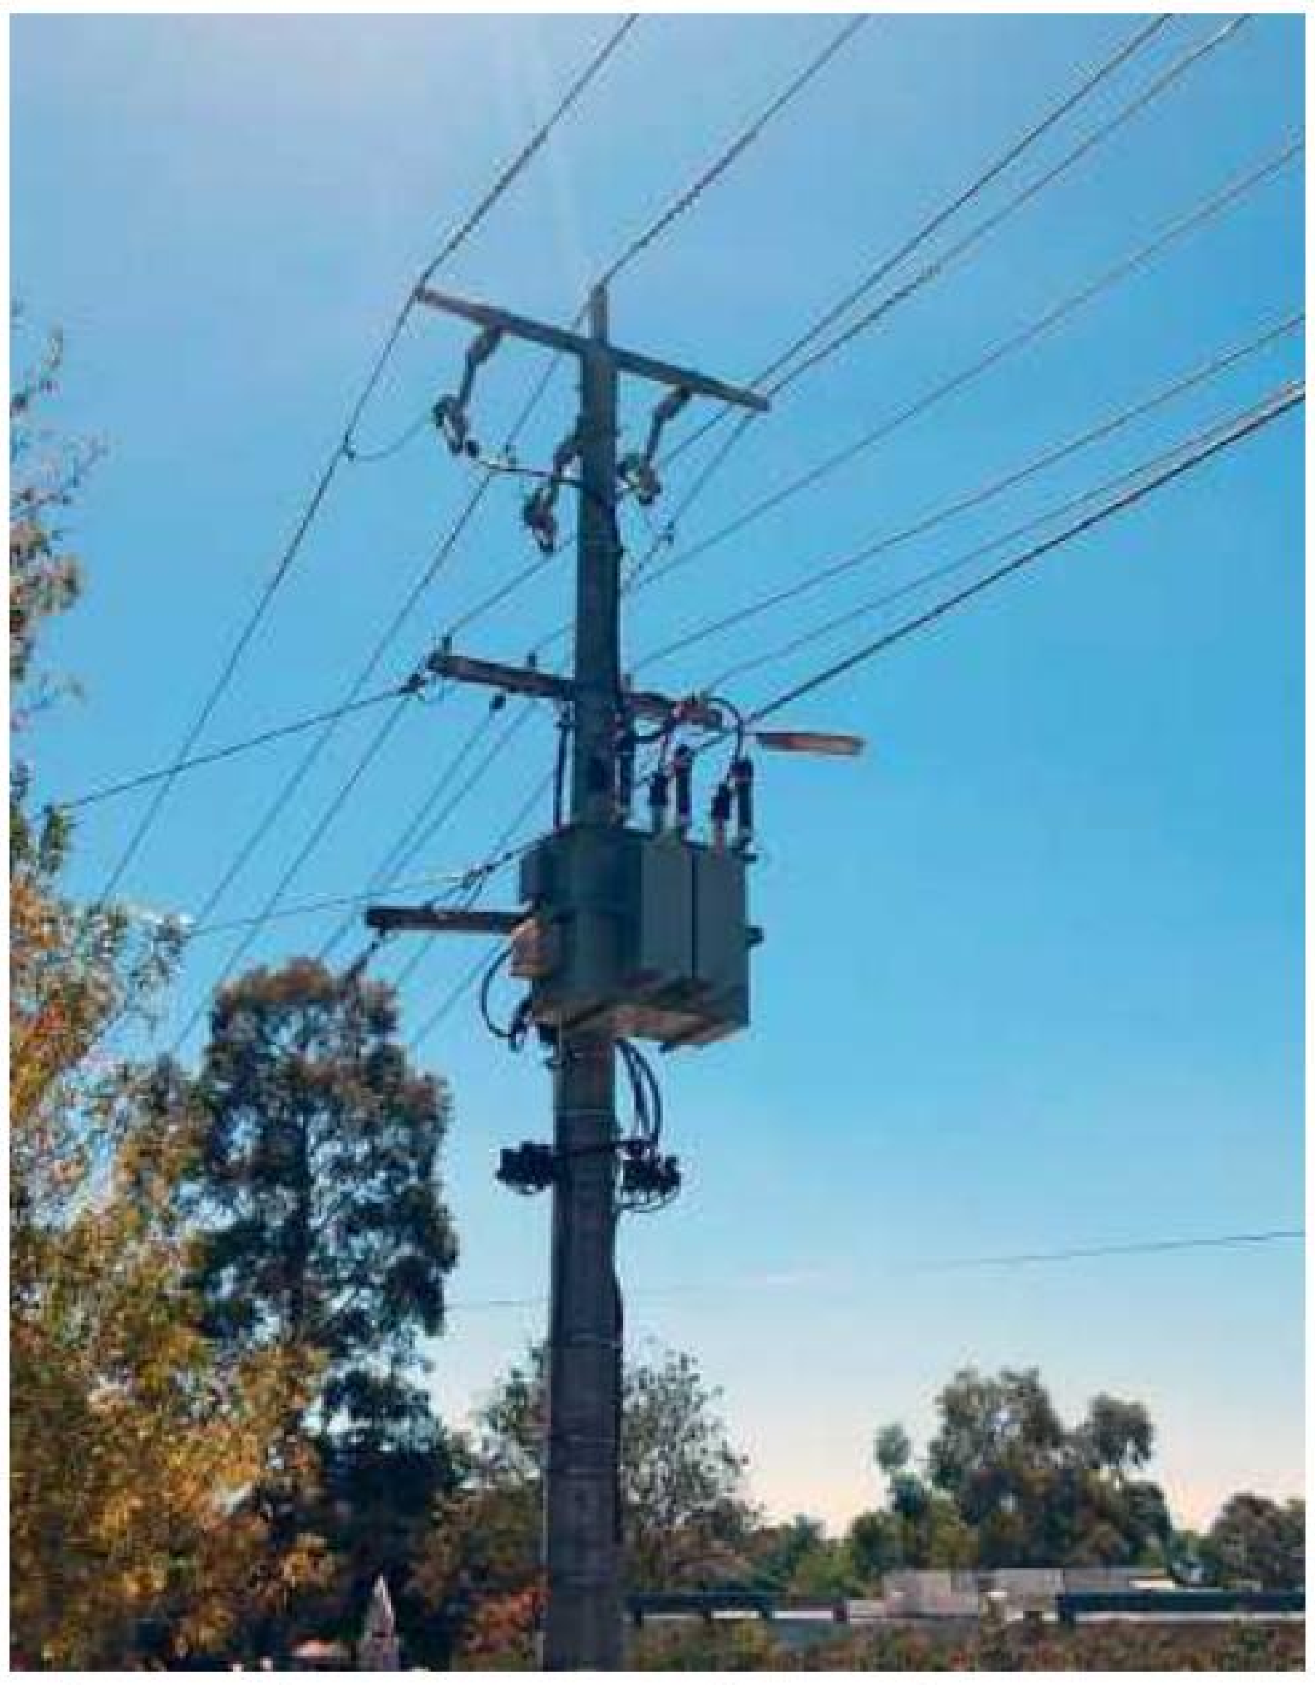

Fig.14: the TN-C-S scheme in Melbourne, Australia. The distribution transformer feeds the four horizontal conductors directly above it. The second conductor from the left is the combined Neutral/Earth (PEN), while the other three conductors are the Line phases, 400V Line-to-Line or 230V Line-to-Neutral. An Earth stake (out of shot) connects to the conductor running up the left of the pole, partially covered by a white conduit. The transformer is fed on the primary side by the three conductors at the top with 22kV between phases.

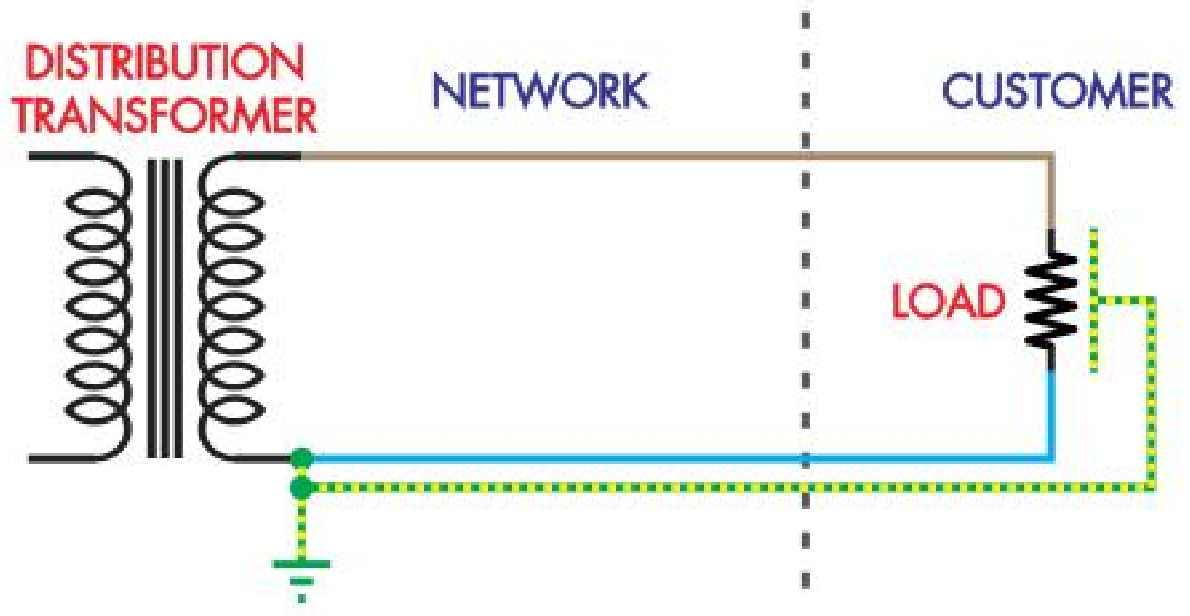

Fig.8: the TN-S Earthing scheme.

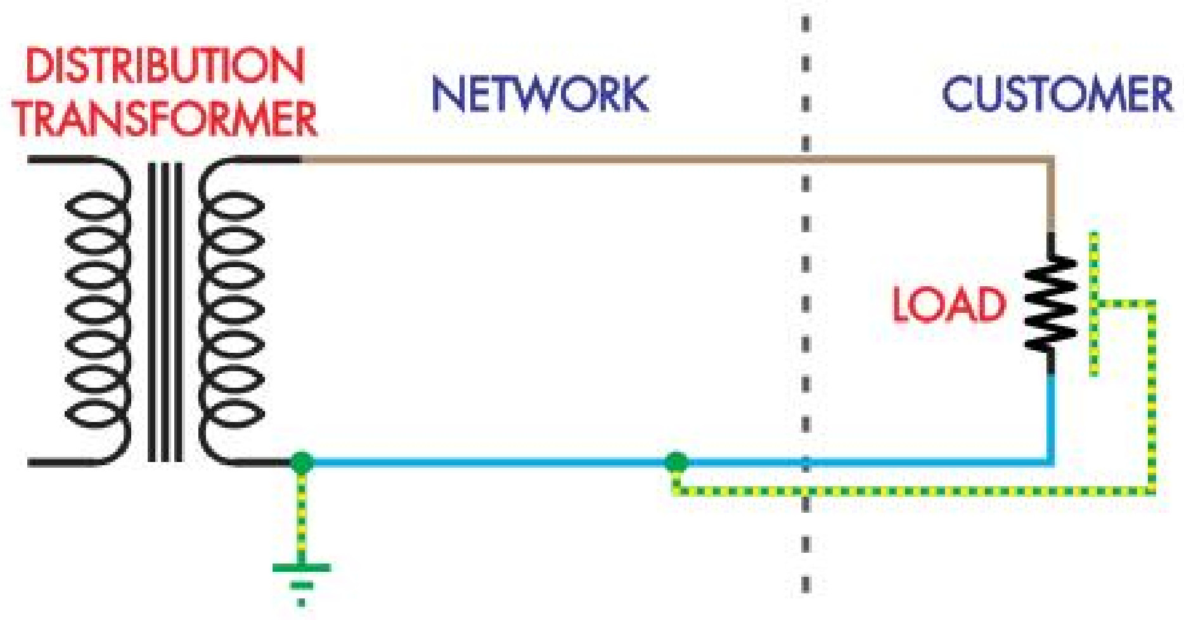

Fig.9: the TN-C-S Earthing scheme.

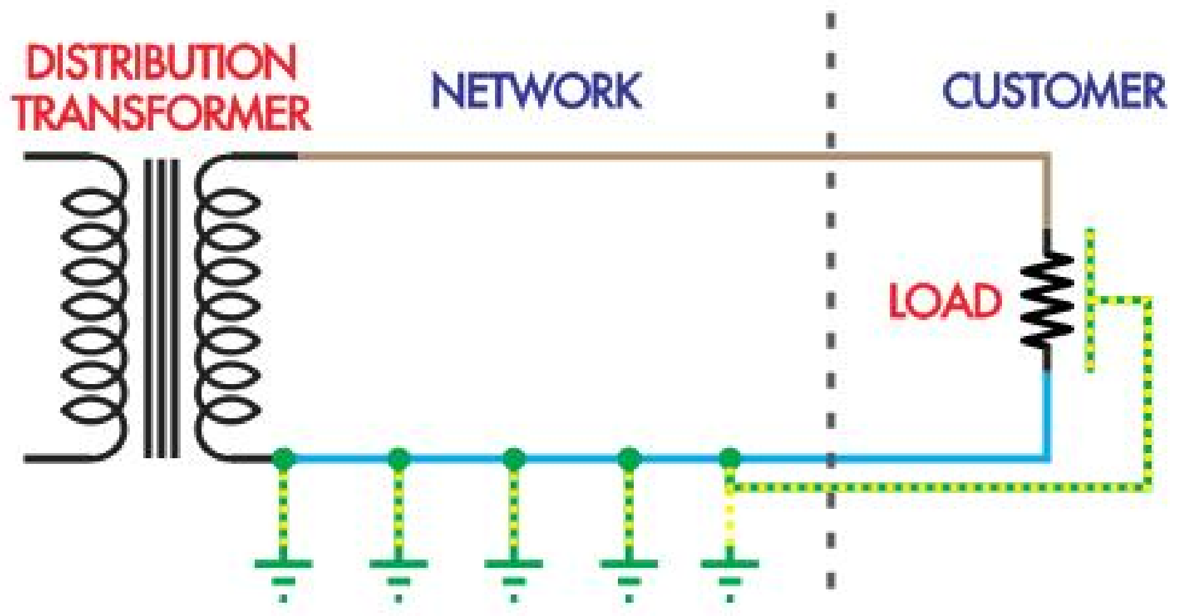

Fig.10: the MEN Earthing scheme.

By Max the Magnificent

WEIRD & WONDERFUL ARDUINO PROJECTS

Part 19: Marble mayhem and tilting tomfoolery

It never ceases to amaze me how often I fall into the trap of thinking something is going to be easy-peasy lemon squeezy to explain, only to discover it quickly transmogrifies itself into a case of stressed, depressed lemon zest.

You only discover what you don't know when you try to explain something you think you know to someone else; several well-known quotes capture this concept. For example, "If you can't explain it simply, you don't understand it" is commonly attributed to Albert Einstein. Likewise, Richard Feynman once observed, "You don't really understand something until you try to teach it".

Sad to relate, this is the position I currently find myself in. As you may recall, last month we said we would add an inertial measurement unit (IMU) to our retro games console.

Although we're exploring all this in the context of our little console, don't be fooled into thinking it's a one-trick pony.

The ability to detect tilt, orientation, and motion turns up in all sorts of applications, from balancing robots and camera stabilisers to gesture-controlled gadgets and industrial monitoring systems. So, while our immediate goals may be modest, the underlying techniques are anything but frivolous.

We can use our IMU to detect whether we are tilting the console left or right, and/or tilting it forward or backward. We will then be perfectly poised to employ this information to control whatever game we happen to be playing at the time (assuming such control is appropriate for the game in question).

This all hinges on what we mean by "forward", "backward", "left" and "right." It would be nice if we were all singing from the same hymn sheet. It also depends on how the IMU interprets those directions.

No, your left, stupid!

Our first task is to define three axes that are perpendicular (at right angles) to each other, as depicted in Fig.1(a). We start by assigning names to them.

Our lives are made easier if we assume we are working in the presence of a gravitational field, because this allows us to immediately define an up–down axis, with the "down" direction pointing toward the gravitational source and the "up" direction pointing away from it.

If we are talking about an object like a boat, we can also say it has a front and a back. In this case, terms like bow (front), stern (back), port (left when looking toward the front), and starboard (right when looking toward the front) make perfect sense because they are defined relative to the vessel itself.

If we are talking about something else, like a landscape, there is no inherent front or back. Those notions depend entirely on the observer. In such cases, we must impose an external frame of reference.

Once we've defined our up–down and front–back axes, we can establish our left–right axis. In this case, "left" lies to the left of the up–down axis when we are looking toward the front (the assignment of "right" is left as an exercise for the reader).

It's common to associate the letters X, Y and Z with these axes, but which letter should go with which axis? Sadly, there's no universally agreed answer. Different disciplines have evolved their own conventions over time. In mathematics, X is typically left–right, Y is up–down, and Z represents depth (into and out of the page or screen, for example).

In robotics, X often points forward, Y to the left, and Z upward. Meanwhile, the practitioners of computer graphics seem to delight in variety, with some systems using Y as the vertical axis and others preferring Z—just to keep everyone on their toes.

For reasons that will become apparent in a moment, we will associate X with left–right, Y with front–back and Z with up–down, as depicted in Fig.1(b).

We have associated increasing (+ve) Z values with up and decreasing (–ve) values with down. Similarly, increasing (+ve) X values correspond to the right and decreasing (–ve) values to the left, while increasing (+ve) Y values

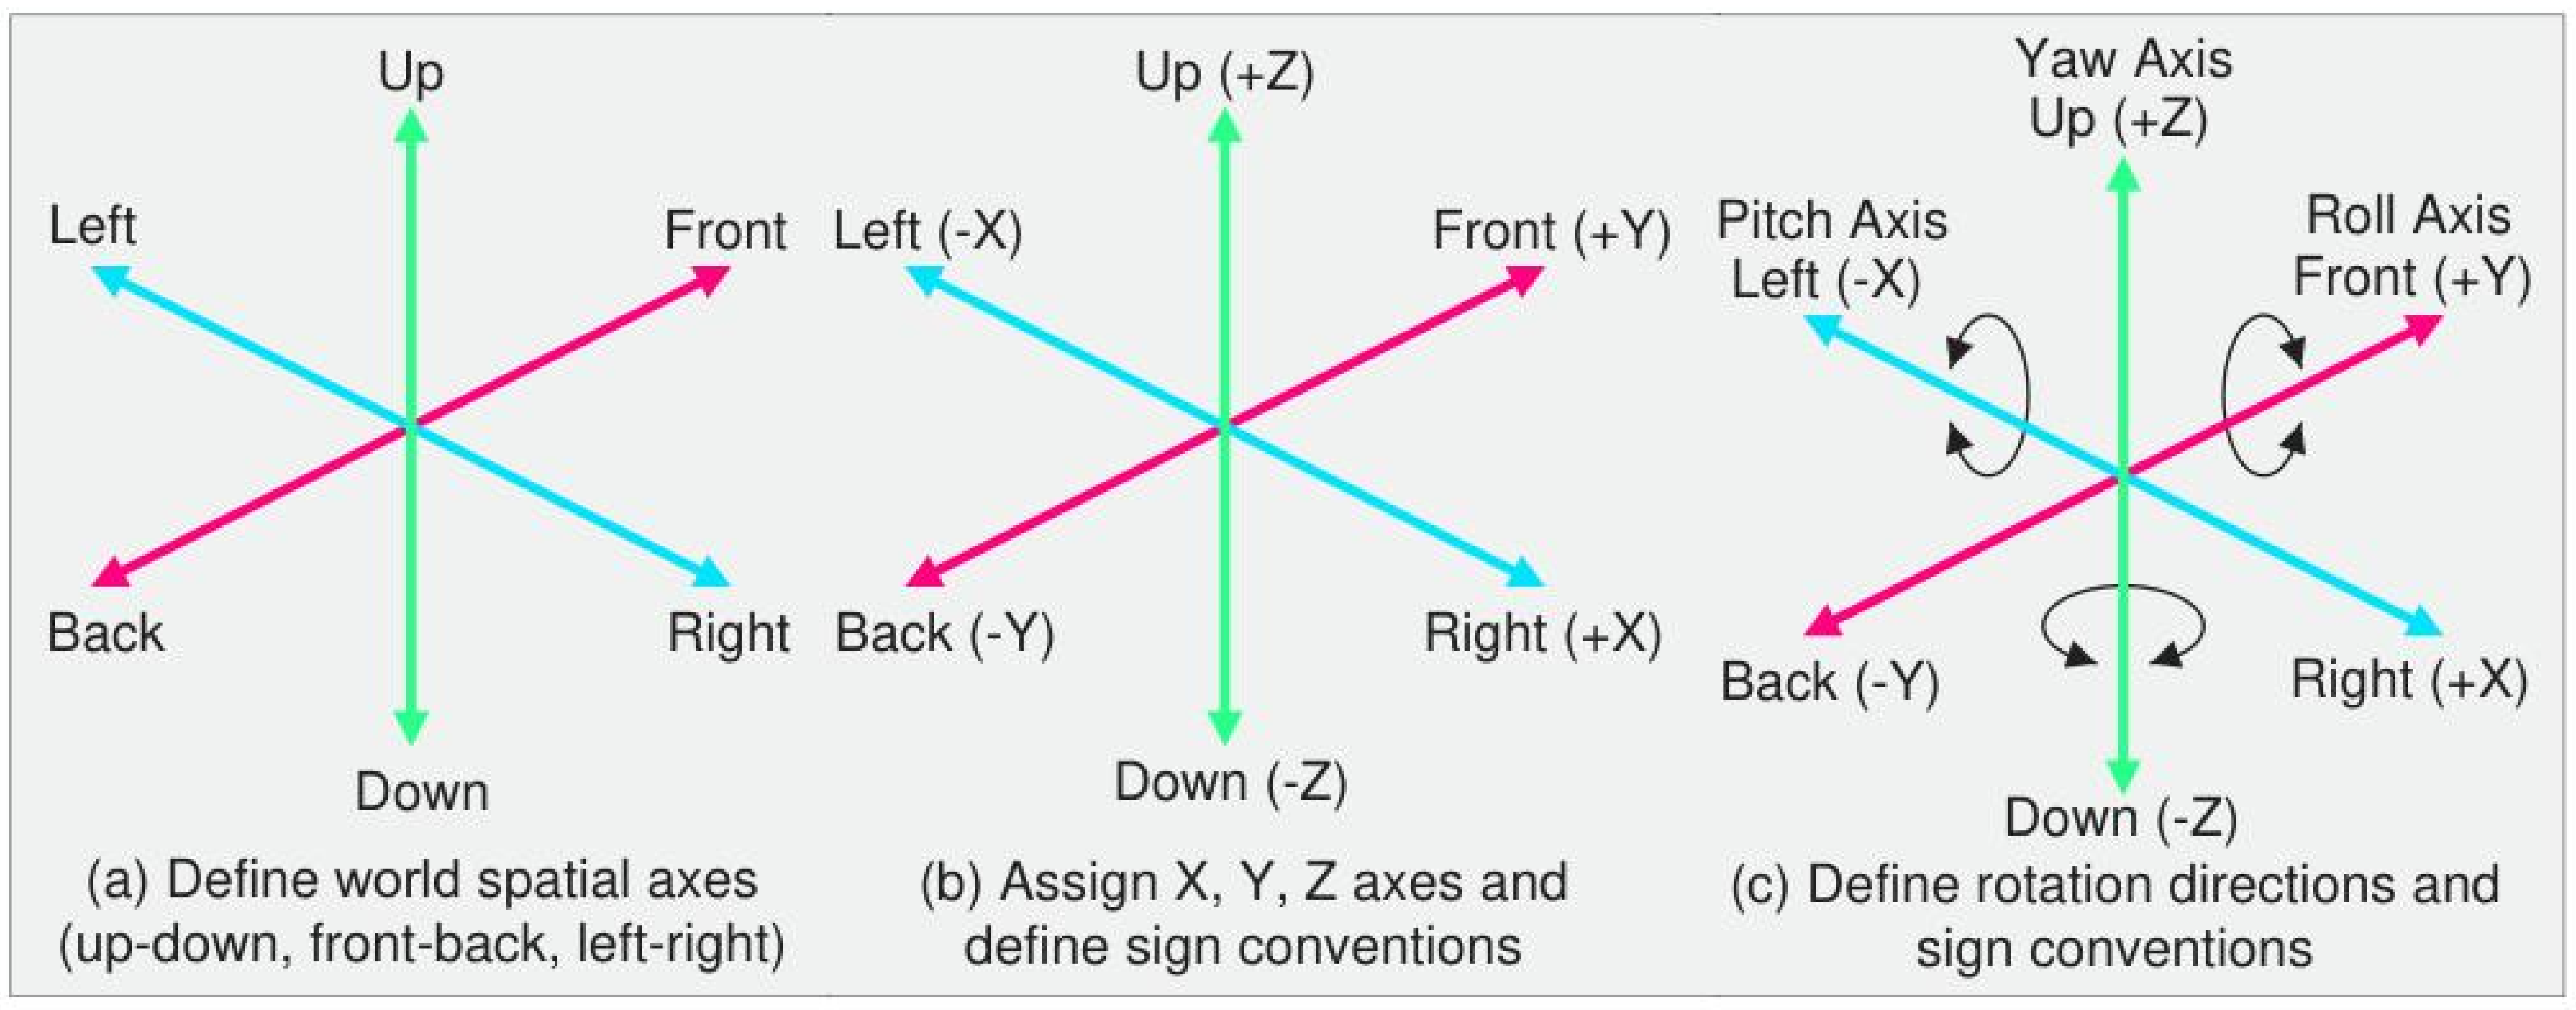

(a) Define world spatial axes (up-down, front-back, left-right)

(b) Assign X, Y, Z axes and define sign conventions

(c) Define rotation directions and sign conventions

Fig.1: definitions and conventions for our IMU's axes and rotations.

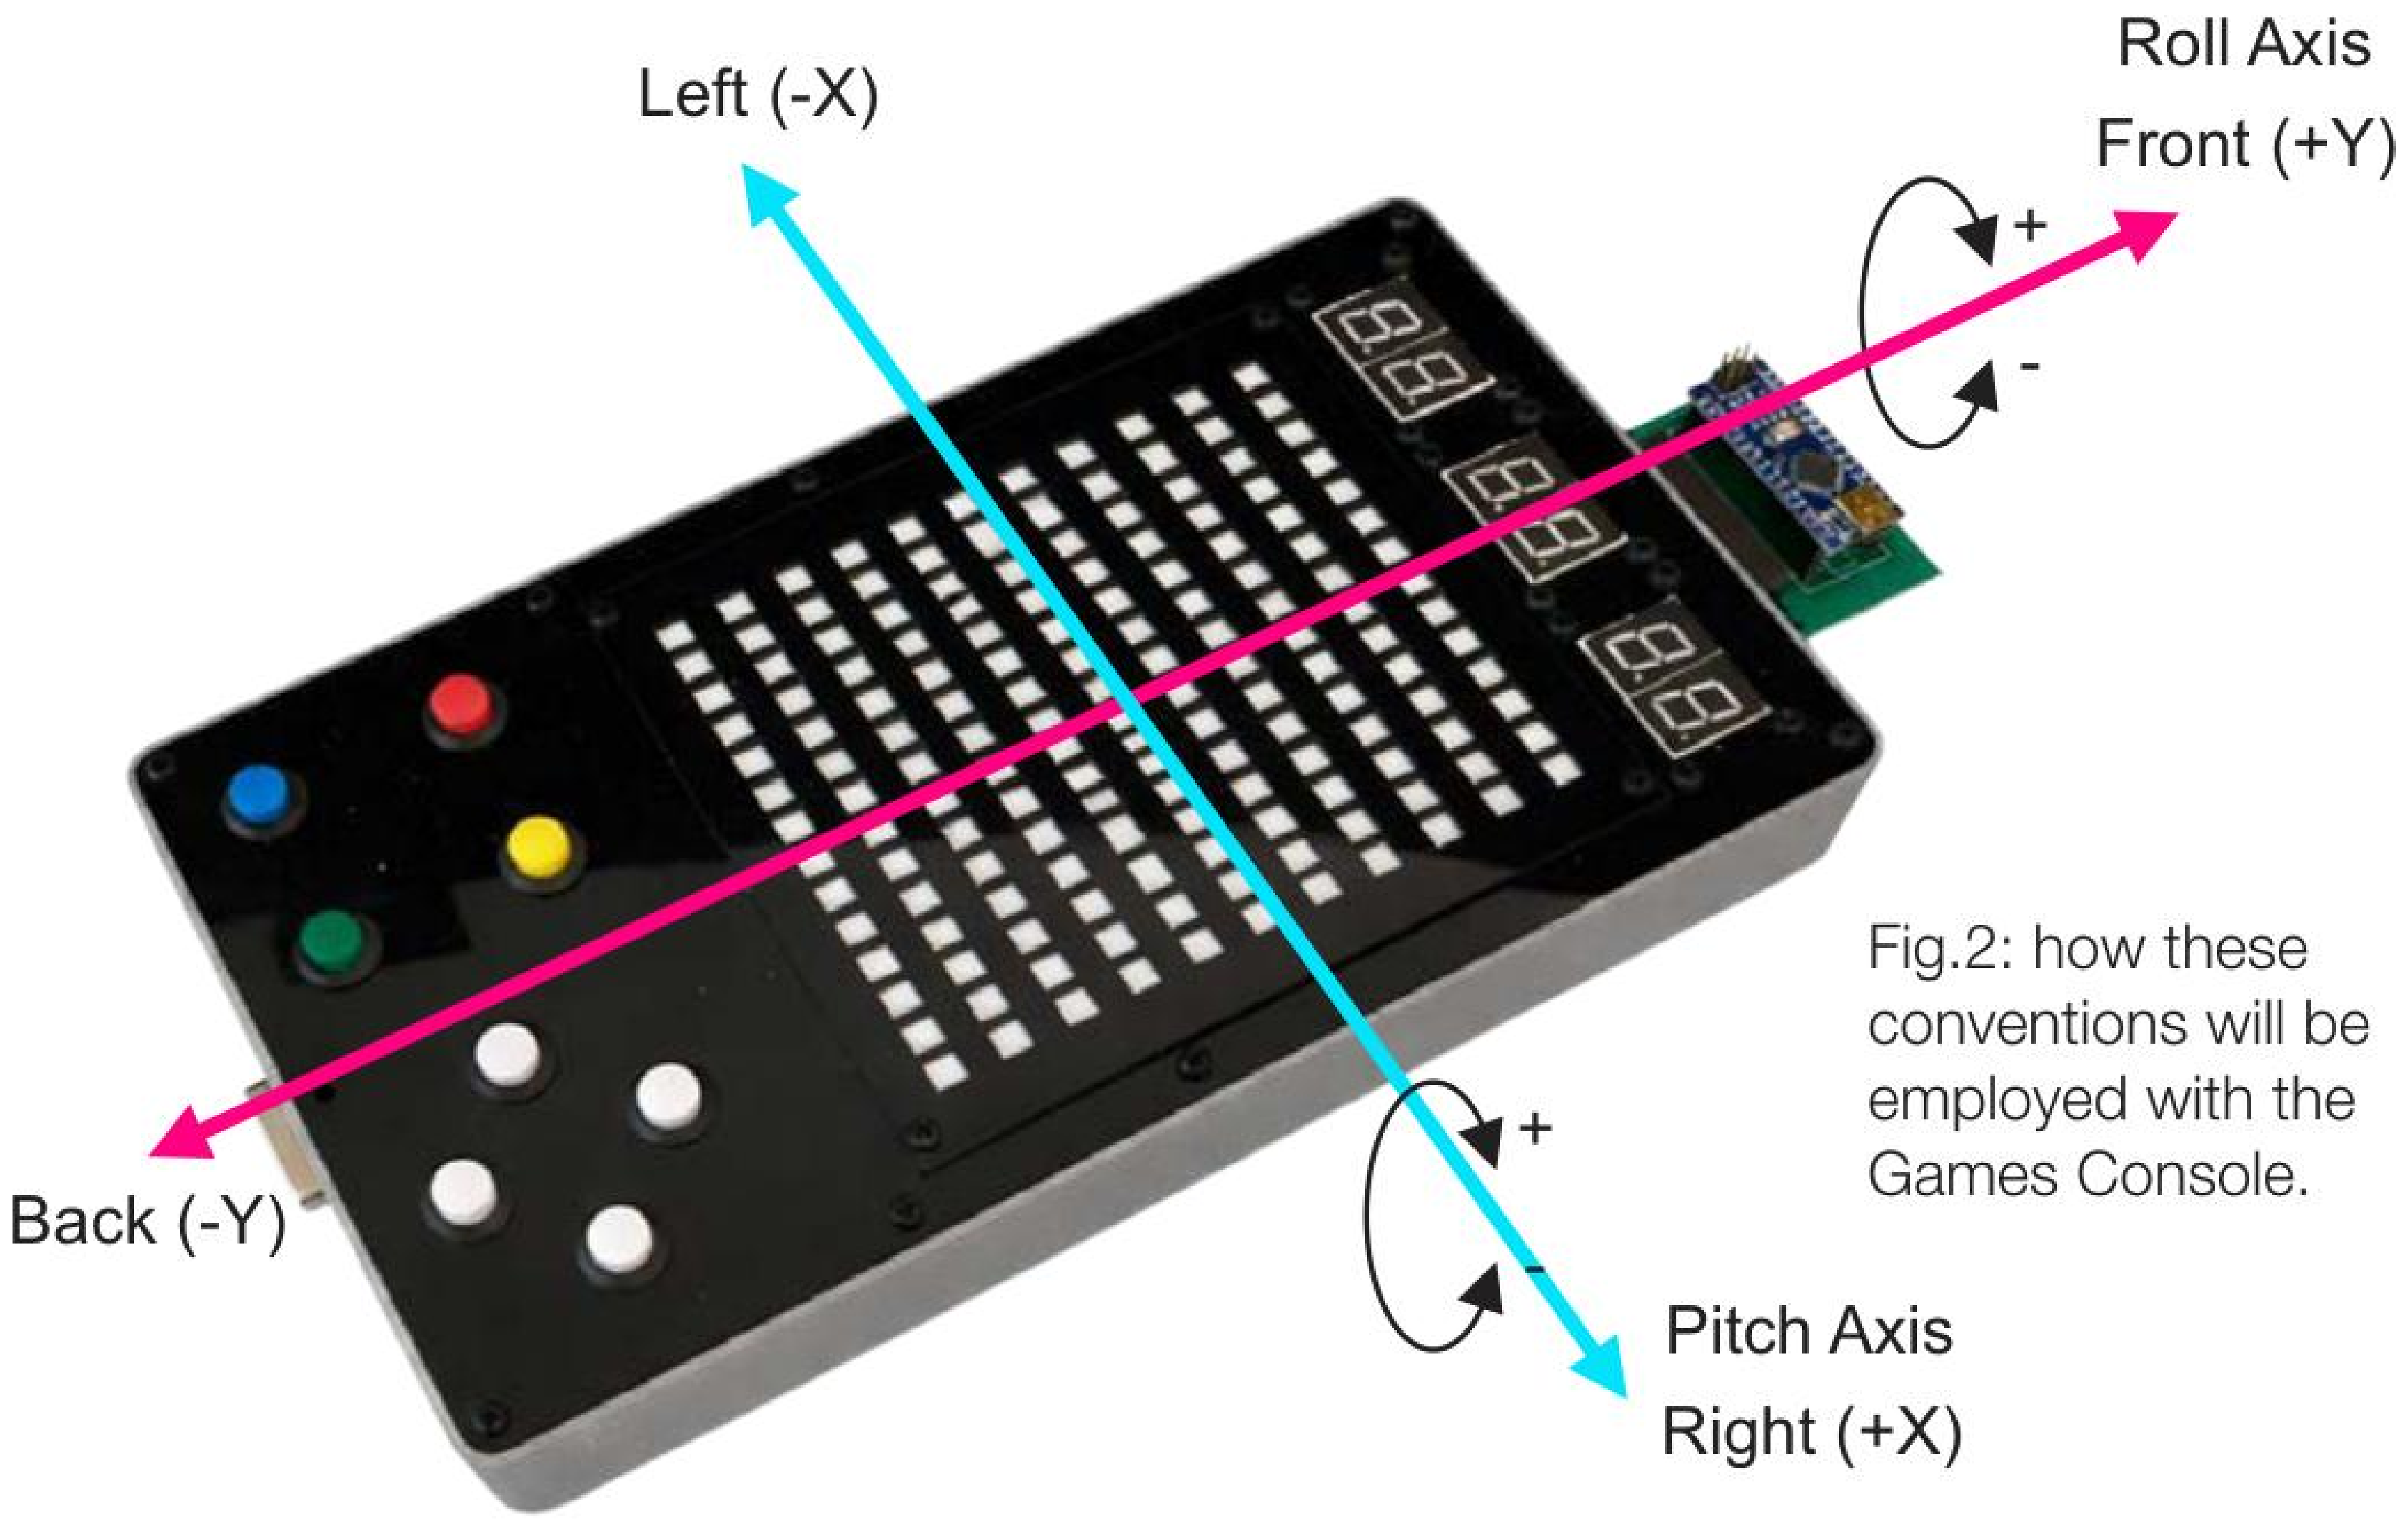

Fig.2: how these conventions will be employed with the Games Console.

correspond to the front and decreasing (–v) values to the back.

These assignments are not universal; different domains have adopted different conventions. In many cases, they have simply chosen the convention that works best for the task at hand, which is exactly what we are doing here.

The terms ‘roll’, ‘pitch’ and ‘yaw’ refer to rotation around the front-back, left-right, and up-down axes, respectively, as illustrated in Fig.1(c). Once again, the signs of these rotations (is clockwise +ve or -ve?) depend on the domain in which we are working. Just to increase the fun and frivolity, we also need to define what we mean by ‘clockwise’ (ie, what frame of reference we’re using).

And then things get complicated

To cut a long story short (which is the opposite of how I usually like to do things), the conventions we will use for our retro games console are shown in Fig.2.

Previously, we established the use of the X and Y axes in the context of our 14 × 10 tricolor light-emitting diode (LED) array. In this case, X spans the width of the array, while Y represents its height. This is why we have chosen to associate X and Y with the left–right and front–back axes, respectively, in the context of our console.

Similarly, the way we reference the array, with the origin at (0, 0) in the lower left-hand corner and (13, 9) in the upper right-hand corner, explains why we have decided to associate positive (+ve) values with front and right, and negative (–ve) values with back and left.

This leaves the Z-axis, which points out of the console and into the third dimension (always a handy place to hide things). We won’t be worrying about rotations about the Z-axis in the context of our console for the foreseeable future, so we’ll largely omit it from our discussions.

When it comes to tilting the console, a tilt to the right (left side up, right side down) corresponds to a roll (rotation about the Y-axis). In our case, we shall define clockwise rotation (when viewed from the back looking toward the front along this axis) as positive because we have to pick something, and this makes sense to me.

Similarly, a tilt forward (back end up, front end down) corresponds to a pitch; rotation about the X-axis. In this case, we shall again define clockwise rotation as positive, but now we are viewing things from the right looking toward the left along this axis.

I know this is probably making your head hurt (it’s certainly not doing mine any good), but wrapping our brains around this now will make things much easier in the not-too-distant future.

One question you may have is why we say that clockwise is defined when looking from back to front (**–ve** to **+ve**) along the Y-axis, while also saying it is defined when looking from right to left (**+ve** to **–ve**) along the X-axis. I’m glad to see that you’re paying attention. This has to do with how I’m visualising things in my head (which is the best place to visualise things, in my experience).

Suppose we were to place a physical glass marble on top of our console and tilt the console to the right (right side down). In this case, our marble will roll toward the right. Since I’m sitting at the back of the console looking toward the front, I naturally perceive this as a clockwise rotation.

If we tilt the console to the right, we will create our code such that a virtual marble ‘rolls’ to the right by keeping the value of Y constant and increasing the value of X. Since we’re making X more positive, it intuitively makes sense (at least, to me) that we should consider this clockwise rotation to be positive.

If we say that a clockwise rotation is positive when we perform a roll about the Y-axis, it makes sense to adopt a consistent convention when we perform a pitch about the X-axis. If we tilt the console forward (front end down), we will make our virtual marble roll toward the front by keeping the value of X constant and increasing the value of Y.

Since we are increasing the value of Y, and since we are associating this motion with a positive rotation, then it follows that this motion must appear as a positive (clockwise) rotation when viewed along the X-axis.

For this to appear clockwise, we must be looking along the X-axis from right to left (that is, from +X toward –X). If we were to look in the opposite direction, the same motion would appear anti-clockwise and our sign convention would be broken.

And people wonder why I drink.

Reconciling multiple realities

Let’s assume we are going to use an IMU that boasts a 3-axis gyroscope (to measure rotation), flaunts a 3-axis accelerometer (to sense gravity and acceleration) and sports a 3-axis magnetometer (to determine orientation).

Each type of sensor has its own set of concerns. Gyroscopes drift over time (bias integration error), accelerometers are noisy and get confused by motion vs gravity, and magnetometers are easily disturbed by nearby magnets, metal and electronics.

As a result, we are faced with multiple ‘truths’ (“The truth is rarely pure and never simple” — Oscar Wilde). The gyroscope says, “I’m rotating smoothly” (while quietly drifting over time), the accelerometer says, “gravity is over here” (except when you’re accelerating), while the magnetometer insists that “north is… well… somewhere in that general direction”.

By combining data from the gyroscopes, accelerometers, and magnetometers, and continuously cross-checking them against one another, advanced sensor fusion algorithms can compensate for the weaknesses of each sensor and provide a much more stable and reliable estimate of orientation and motion than any one sensor could achieve on its own.

You can perform sensor fusion on your own, assuming you’re happy to juggle Euler angles, wrestle with quaternions, wrangle rotation matrices, normalise vectors, transform coordinates and face the ever-present spectre of numerical drift. Alternatively—and this is the option I prefer—you can get someone else to do it for you.

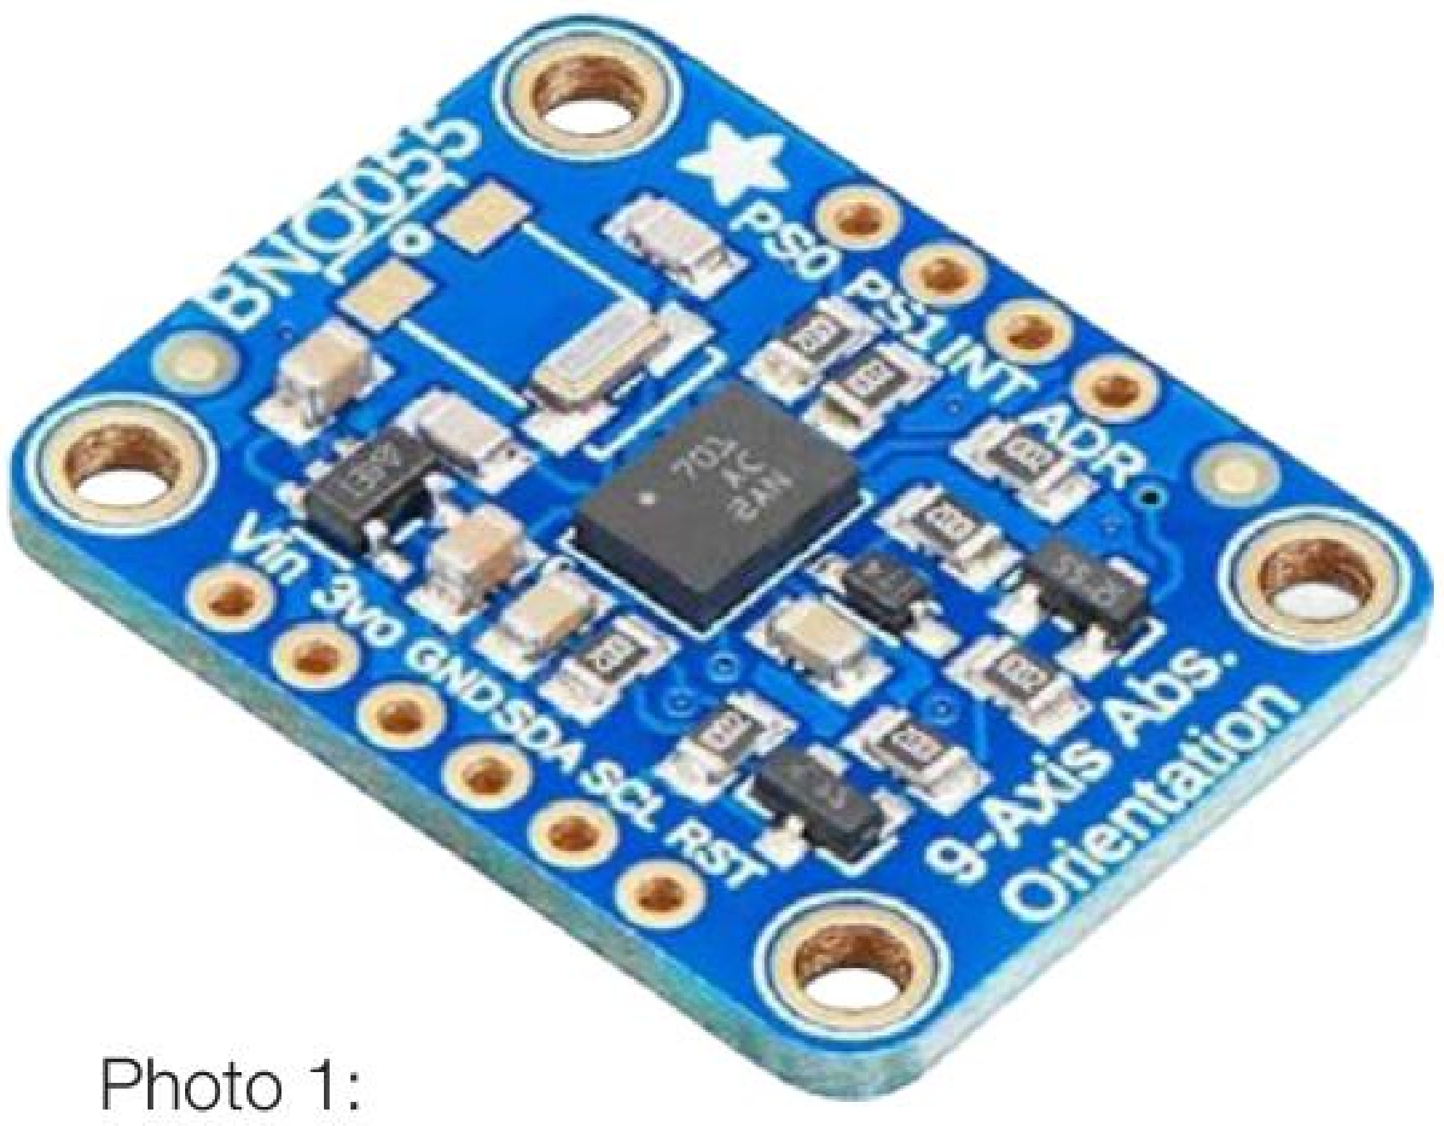

Photo 1:
Adafruit's BNO055 breakout board.

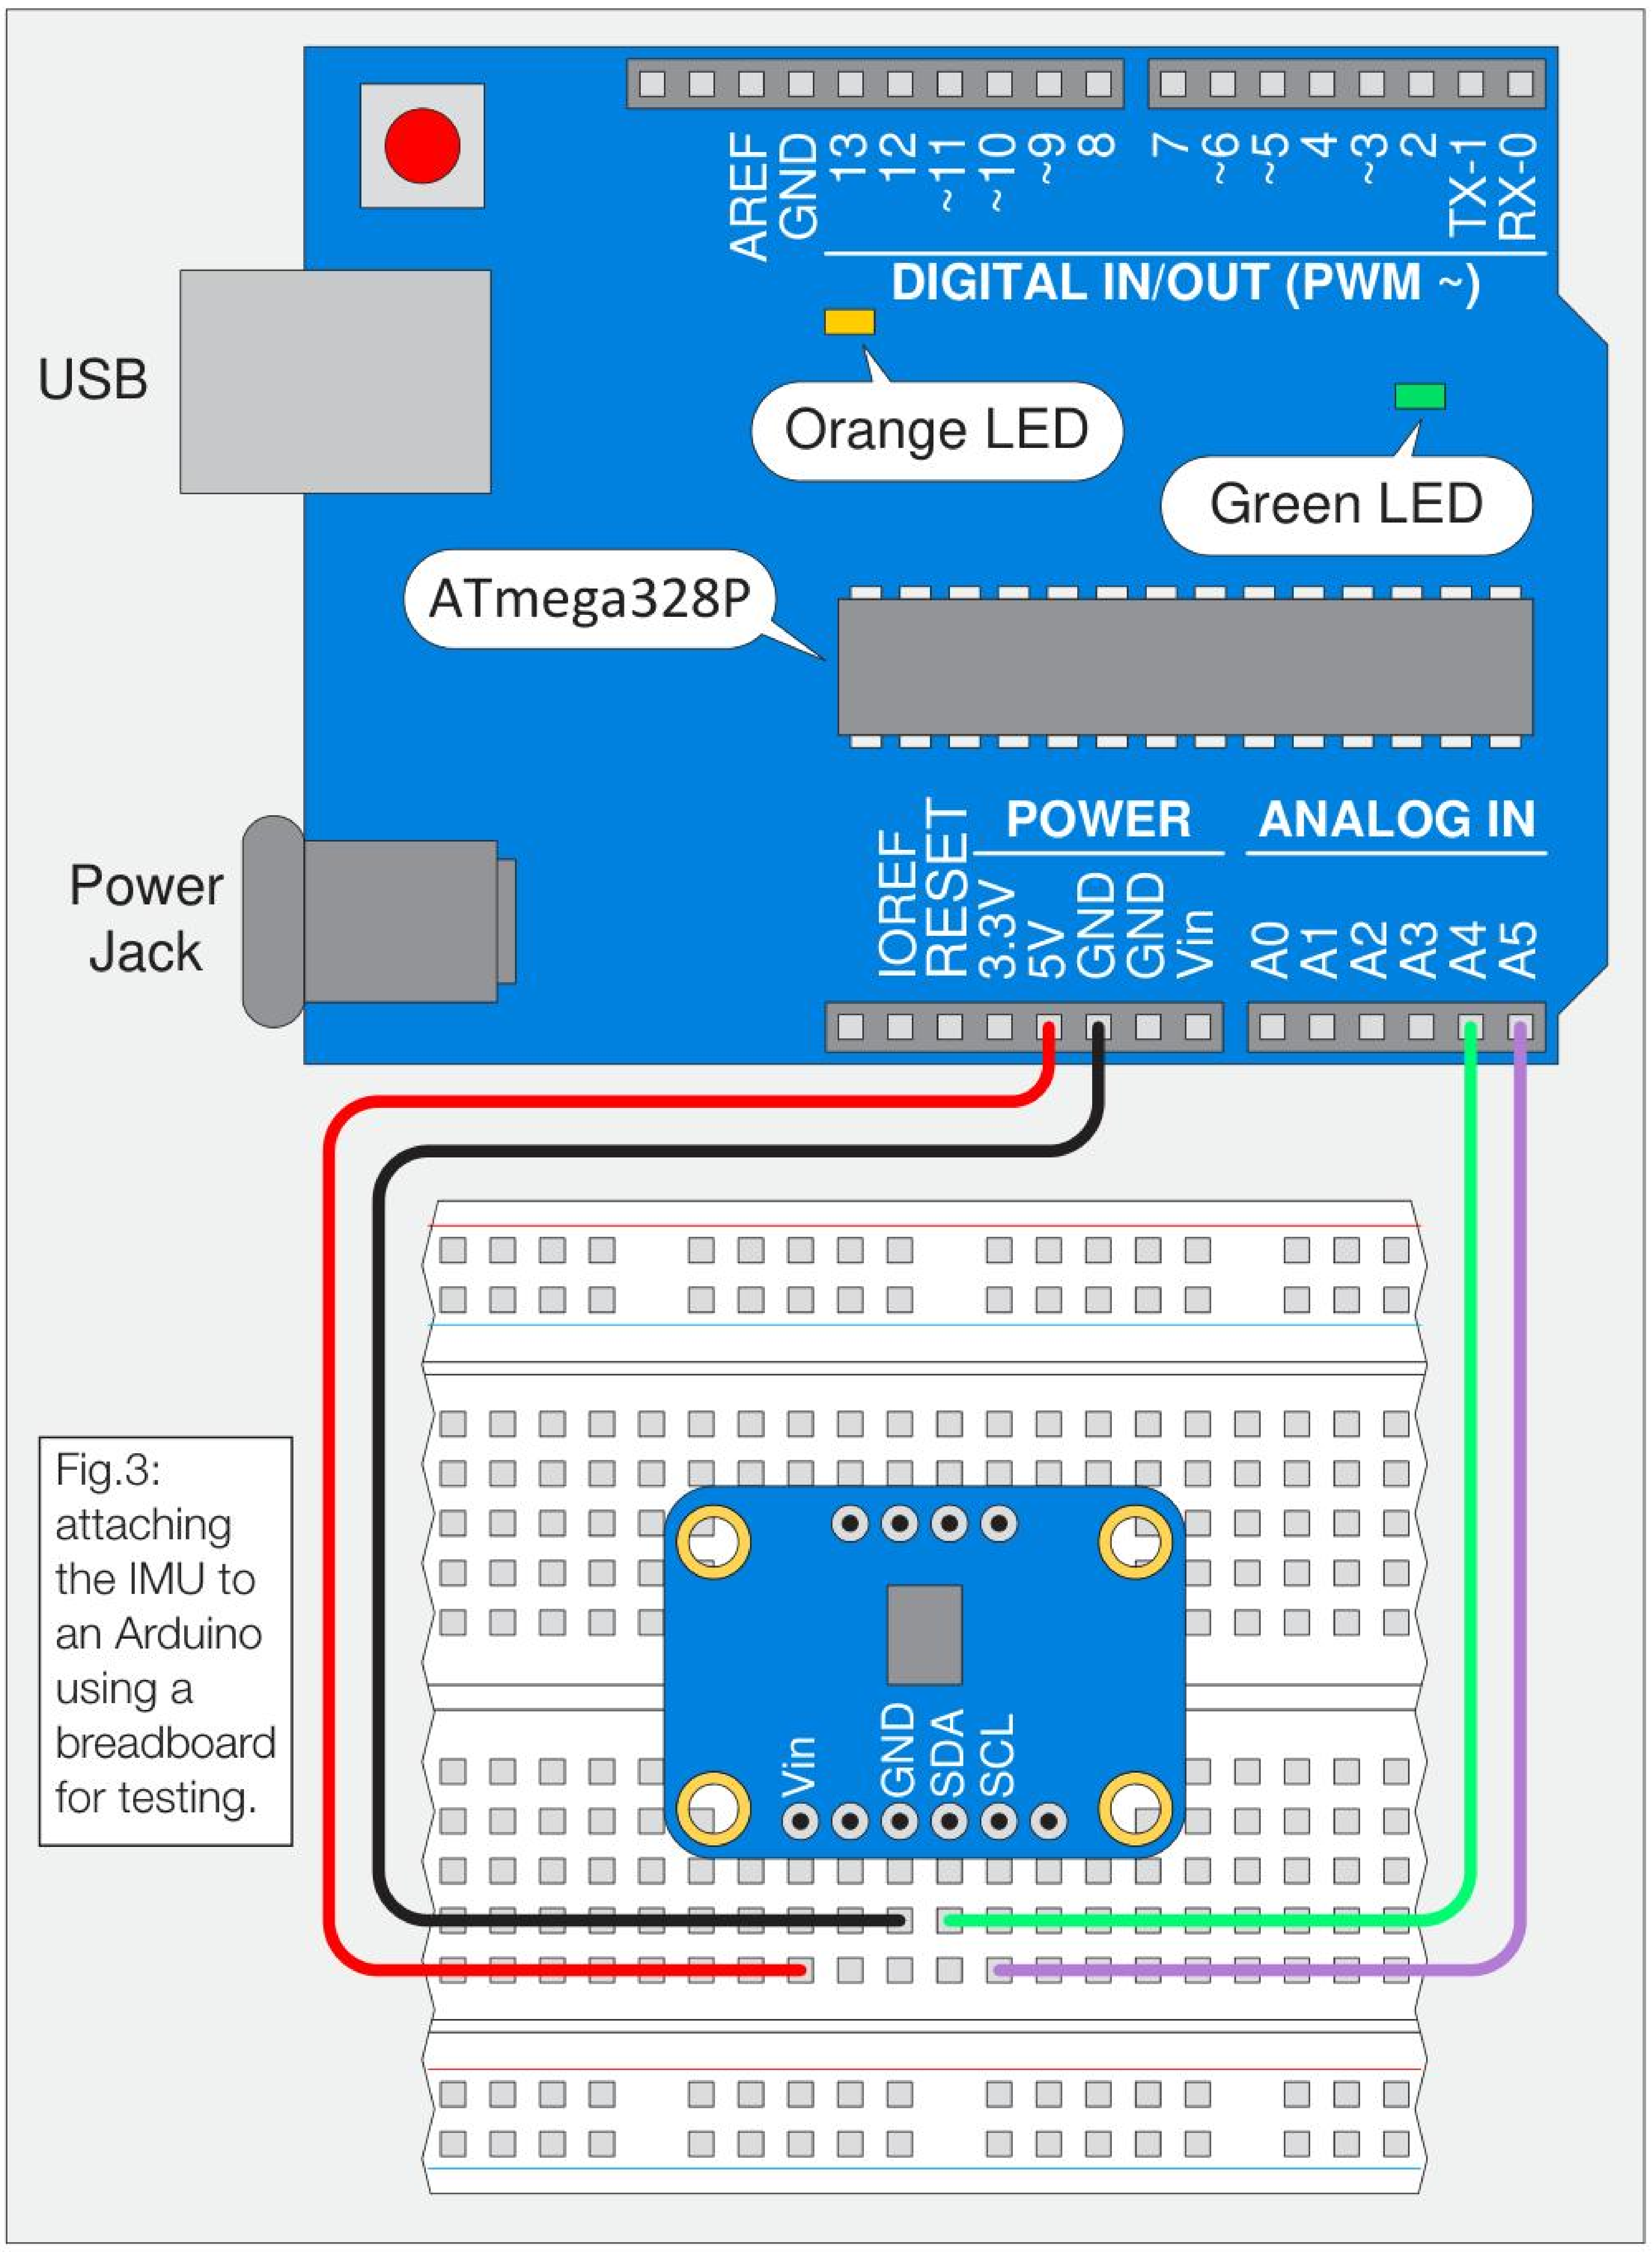

Fig.3: attaching the IMU to an Arduino using a breadboard for testing.

Meet our IMU

As we noted last month, there are various cheap-and-cheerful IMUs on the market. Although it's not the bargain-basement option, I do have a soft spot for the 9DOF Absolute Orientation IMU Fusion Breakout board from Adafruit (**www.amazon.co.uk/dp/B017PEIGIG**) – see Photo 1.

There are three reasons why I like this little beauty, which is based on the Bosch BNO055 micro-electromechanical system (MEMS) sensor:

1. The BNO055 includes a 32-bit microcontroller unit (MCU), which the nice guys and gals at Bosch have already programmed to perform sensor fusion.
2. The chaps and chapesses at Adafruit have done their usual outstanding job creating and providing wiring information, example code, and luscious libraries to get us up and running (**www.adafruit.com/product/2472**).
3. I already have one of these little scamps lurking in the depths of my treasure chest of "Bits and pieces that will undoubtedly come in useful one day if I only wait long enough".

Can you guess how the BNO055's creators decided to handle all the X-Y-Z and pitch-roll-yaw stuff? I don't have a clue. We could look at the data sheet, but what would be the fun in that? (I also have a sneaking suspicion that attempting to work things out from the datasheet would cause my brain to leak out of my ears.)

After all, how the IMU's X-Y-Z axes ultimately map onto the console's X-Y-Z axes will depend on how we end up physically mounting the IMU in (or on) the console. So, on this basis, I think our best bet is to adopt the time-honoured engineering approach of "suck it and see".

Setting up the hardware

Before we integrate the BNO055 breakout board into our console, we'll begin by experimenting with it in isolation. This will let us get up to speed with its capabilities, understand its quirks, and iron out any problems in a simple setup before introducing any added complexity to the full system.

Since my IMU already had its header pins soldered in, I plugged it into a breadboard and then connected the breadboard to a spare Arduino Uno I had lying around, as in Fig.3.

The BNO055 sensor communicates with the outside world using a simple two-wire serial bus interface called I^{2}C. One wire carries the clock signal (called **SCL**), which keeps everything in step, while the other carries the data (called **SDA**).

One device—usually a microcontroller like an Arduino Uno—acts as the boss and controls the conversation with any other devices ("targets") on the bus. These additional devices can include things like sensors, displays and memory chips. On both the Uno and Nano, analog pins **A4** and **A5** serve as the I^{2}C **SDA** and **SCL** signals, respectively.

Each device has its own address, rather like a house number. When the Arduino wants to talk to a particular device, it first sends that device's address. All the devices listen, but only the one with the matching address responds. Thus, many devices can share the same two wires, which saves a lot of pins and wiring. That's why I^{2}C is so popular in small electronic systems.

It's important to note that the I^{2}C bus requires a pull-up resistor on each of its wires, but these are included on the BNO055 breakout board, so that's one less thing for us to worry about (phew!).

Setting up the software

On the software side, the first thing we need to do is install the libraries required to run our BNO055 sensor (if you use a different IMU, you'll need to load its associated libraries).

Launch the Arduino IDE, then use the **Tools → Manage Libraries** menu commands to access the Library Manager. Search for "Adafruit BusIO" and install it. This is a low-level helper library that many Adafruit sensor libraries depend on.

```
 1  #include <Wire.h>
 2  #include <Adafruit_Sensor.h>
 3  #include <Adafruit_BNO055.h>
 4  #include <utility/imumaths.h>
 5
 6  Adafruit_BNO055 IMU = Adafruit_BNO055(55);
 7
 8  void setup()
 9  {
10    Serial.begin(115200);
11    while (!Serial) {}
12
13    Serial.println("BNO055 orientation test");
14    Serial.println("Move the breadboard, watch which values change.");
15    Serial.println();
16
17    if (!IMU.begin()) {
18      Serial.println("Ooops, no BNO055 detected ... check wiring!");
19      while (1) {}
20    }
21
22    delay(1000);
23    IMU.setExtCrystalUse(true);
24  }
25
26  void loop()
27  {
28    sensors_event_t event;
29    IMU.getEvent(&event);
30
31    Serial.print("IMU.x = ");
32    Serial.print(event.orientation.x, 2);
33    Serial.print("\tIMU.y = ");
34    Serial.print(event.orientation.y, 2);
35    Serial.print("\tIMU.z = ");
36    Serial.println(event.orientation.z, 2);
37
38    delay(100);
39  }
```

Listing 1(a): our first IMU test program.

Next, search for "Adafruit Unified Sensor" and install that. This provides the common sensor framework that the BNO055 library will use.

Finally, search for "Adafruit BNO055" and install that library last. This is the actual sensor driver for the BNO055.

At this point, I suggest you exit and relaunch the IDE because the Arduino IDE doesn't always immediately recognise newly installed libraries. Restarting saves you from mysterious "file not found" errors later on. You're welcome.

Coding our first test

The code for our first test is essentially a copy of the example code on Adafruit's website. This is a nice, simple "Is the sensor alive and what happens when I tilt it?" test program. You can find the version we're using in the file named **CB-jul26-code-01.txt** (as usual, all files mentioned in this column are in the July 2026 download package from the PE website at https://pemag.au/link/acbx).

This code is shown in Listing 1(a). We're using the convention that the listing number (1 in this case) corresponds to the numerical part of the matching code file name ("01" here).

Since much of this is new to us, it's worth looking at it in detail. We start by including several libraries. On line 1, we include **<Wire.h>**, which is the Arduino's built-in I²C library.

On line 2, we include **<Adafruit_Sensor.h>**, which is Adafruit's common sensor framework. This provides standard data structures and helper functions used by many sensors.

On line 3, we include **<Adafruit_BNO055.h>**, which is the actual library for the BNO055. This provides us with the sensor-specific functions we'll use later, like **IMU.begin()** and **IMU.getEvent()**.

On line 4, we include **<utility/imumaths.h>**, which provides support for the vector and math types used by the BNO055 library.

On line 6, we create our IMU object, which we've cleverly called "IMU" (don't worry about the number 55, which is just an identifier used internally by Adafruit's framework).

On line 10, in the **setup()** function, we set the serial port running at 115,200 baud (think "bits per second"). This is much higher than the value of 9600 we usually use because we're going to stream multiple sensor readings several times per second. The higher data rate keeps the Serial Monitor updated smoothly and helps avoid bottlenecks, such as data queuing faster than it can be displayed.

The statement on line 11 waits until the serial connection is ready for use. This matters more on some boards than on an Uno, but it's harmless and good practice to include it.

Lines 13, 14 and 15 are used to print a heading to the Serial Monitor. This is the message we'll see when we open the monitor.

Line 17 attempts to start the BNO055. The library checks whether the sensor is connected and responds over the I²C bus. If the initialisation fails, we print an error message (line 18) and then pause forever (line 19). That way, we won't confuse ourselves by displaying nonsense readings.

Assuming everything has passed thus far, we pause for a second (line 22) to give the sensor time to settle. Next, we tell the sensor to switch over to using the crystal oscillator on the breakout board as its clock (line 23). This improves timing stability and makes the sensor's readings more reliable.

Now let's turn our attention to the **loop()** function. This runs repeatedly forever. Each time around the loop, we ask the sensor for its orientation, print its responses, and wait for 1/10th of a second (100ms).

On line 28, we create a variable called **event**. This is of type **sensors_event_t**, which is defined by the BNO055 library. Think of this as a little container that will hold data from the sensor. On line 29, we use **IMU.getEvent()** to fetch the current orientation readings from the sensor and store them in our **event** variable.

On lines 31-36, we print the **orientation.x**, **.y**, and **.z** values stored in our event variable. These are floating-point values, meaning they can include both whole-number and fractional parts. The **\t** portions of these statements insert tabs to make the columns line up nicely, while the **2** values tell the Arduino to print the floating-point numbers with two digits after the decimal point.

Running our first test

Load this program into your Arduino. It will start running automatically. Open the Serial Monitor window and make sure it's set to 115200 baud, matching the value in our code.

In my case, I had everything arranged as shown in Fig.3. I was sitting at my desk with the breadboard directly in front of me. Once the program was up and running, I tilted my breadboard about 45° to the left, then 45° to the right, then 45° forward, and finally 45° backward. The first few lines from my run look like the following (note that I've chopped a lot of stuff out).

```
BNO055 orientation test
Move the breadboard, watch
        which values change.
IMU.x=359.25  IMU.y=-4.25   IMU.z=1.44
IMU.x=358.94  IMU.y=-10.13  IMU.z=0.69
```

These two lines correspond to tilting the breadboard to the left. From this, you might assume that working with these values would be a piece of cake. You'd be mistaken. If you look at all the results from my first run (see the file named **CB-jul26-results-01.txt**), it quickly becomes apparent that nothing makes sense (at least, not to the untrained eye).

Try experimenting with tilting your own breadboard and see if you can work out what's going on.

Part of the problem stems from the fact that the **orientation.x**, **.y**, and **.z** values are presented as Euler angles. These are named after the 18th-century Swiss polymath (and mathematic fanatic) Leonhard Euler. Unfortunately, these values are interrelated in non-obvious ways.

You know how, with ordinary numbers, you can swap the order of the operands in a multiplication and the result remains the same (for example, 6 × 3 = 3 × 6)? This is known as the commutative property. Well, that doesn't apply to matrices.

Although we want to think of Euler angles as three separate rotations, what's really going on is a sequence of matrix multiplications. Since matrix multiplication is not commutative, the order of these rotations matters.

As a result, first tilting to the left and then tilting forward will give a different set of Euler angles than first tilting forward and then tilting to the left. And from here on, things start to get… interesting.

Adding gravity vectors

Although Euler orientation values are useful for all sorts of things, they are not ideally suited for our purposes here—that is, detecting the tilt angles of our console. Fortunately, the BNO055 can also produce a gravity vector with **.x**, **.y**, and **.z** components. We can think of this vector as the 'down' direction expressed in the sensor's coordinate frame, which makes it ideal for 'marble rolling' style behaviour.

Let's keep our current Euler values for curiosity, whilst also adding **grav.x**, **.y**, and **.z** values to the output. The bulk of our program remains the same. All we need to do is modify the **loop()** function, as illustrated in Listing 2(a) from the file named **CB-jul26-code-02.txt**.

On line 30, we declare a variable called **event** as before. On line 31, we use the **getEvent()** function to query the sensor for its orientation and store the results in our **event** variable.

On line 33, we declare a three-element vector variable called **grav**. Later in this line, we use the **getVector()** function to ask the sensor 'which way is down?' and store the results in our **grav** variable.

I know this statement looks like punctuation gone wild; I'm sorry, it's because we've wandered beyond traditional C into the more exotic realms of C++, with its templates, namespaces and object-oriented syntax. All that is beyond the scope of our discussion here, so let's just accept that it works and move on with our lives.

On lines 35-37, we print the **orientation.x**, **.y**, and **.z** values stored in our **event** variable, each with two digits after the decimal point. On lines 38-40, we print the **grav.x()**, **.y()**, and **.z()** values stored in our **grav** variable, again with two digits after the decimal point.

Once again, as soon as my program was up and running, I tilted my breadboard about 45° to the left, then 45° to the right, then 45° forward, and finally 45° backward. The first few lines from this run look like the following:

```
28  void loop()
29  {
30    sensors_event_t event;
31    IMU.getEvent(&event);
32
33    imu::Vector<3> grav = IMU.getVector(Adafruit_BNO055::VECTOR_GRAVITY);
34
35    Serial.print(event.orientation.x, 2);  Serial.print('\t');
36    Serial.print(event.orientation.y, 2);  Serial.print('\t');
37    Serial.print(event.orientation.z, 2);  Serial.print('\t');
38    Serial.print(grav.x(), 2);             Serial.print('\t');
39    Serial.print(grav.y(), 2);             Serial.print('\t');
40    Serial.println(grav.z(), 2);
41
42    delay(100);
43  }
```

Listing 2(a): our second IMU test program.

```
BNO055 tilt test
Tilt left-right and front-back
        and watch what changes.
EulerX EulerY EulerZ GravX GravY GravZ
1.44   -43.06  3.63  -6.69 -0.45  7.14
1.44   -43.06  3.63  -6.69 -0.45  7.14
```

These two lines correspond to tilting the breadboard to the left. As before, I've chopped a lot of stuff out.

Although this may look a little messy, if you examine the results from the full run (see the file named **CB-jul26-results-02.txt**), it becomes much easier to see what's going on. Tilting the board about 45° to the left yields **GravX** values of around -6.9, while tilting it about 45° to the right produces **GravX** values of around +6.4. In both cases, **GravY** remains close to zero.

By comparison, tilting the board about 45° forward returns **GravY** values of around +7.0, while tilting it about 45° backward produces **GravY** values of around -7.4. In both these cases, **GravX** remains close to zero.

All this tells us that the **GravX** and **GravY** values from our sensor correspond to the X and Y axes we've adopted for our console, respectively. Even better, the signs of the **GravX** and **GravY** values map directly onto those we need for our console. What are the odds? (We could, of course, have corrected any mismatches under program control, but it's nice not to have to do so.)

Meanwhile, the **GravZ** values remain positive and relatively large throughout the entire run. From this, we can deduce that **GravZ** represents the vertical component of gravity passing through the board. When the breadboard is lying flat and face-up, **GravZ** is close to +9.8. As we tilt the board by about 45°, some of that gravity is diverted into the X or Y direction, so **GravZ** falls to around +7.

If we were to hold the board vertically, **GravZ** would move toward zero, while flipping the board upside down would cause **GravZ** to go negative.

Converting to angles

We're in a pretty good place at this point. What we're going to do now is use only the X and Y components of the gravity vector from our sensor. We will also limit the effective tilt to ±60° in any direction by clipping the calculated angles in software.

There are two reasons for doing this. First, in practice, we're unlikely to tilt the console much beyond 60° during normal use, so allowing for larger angles doesn't buy us anything. Second, as we approach steeper angles, the response becomes

```
37 void loop()
38 {
39   imu::Vector<3> grav = IMU.getVector(Adafruit_BNO055::VECTOR_GRAVITY);
40
41   // Normalize and protect against tiny over/under-shoots
42   float gx = clampFloat(grav.x() / G, -1.0, 1.0);
43   float gy = clampFloat(grav.y() / G, -1.0, 1.0);
44
45   float leftRightDegF = asin(gx) * 180.0 / PI;
46   float frontBackDegF = asin(gy) * 180.0 / PI;
47
48   int leftRightDeg = (int)round(leftRightDegF);
49   int frontBackDeg = (int)round(frontBackDegF);
50
51   // Optional clipping for display / gameplay
52   leftRightDeg = constrain(leftRightDeg, -MAX_TILT, MAX_TILT);
53   frontBackDeg = constrain(frontBackDeg, -MAX_TILT, MAX_TILT);
54
55   Serial.print("LR = ");   Serial.print  (leftRightDeg);
56   Serial.print("  FB = "); Serial.println(frontBackDeg);
57
58   delay(100);
59 }
```

Listing 3(a): our third IMU test program.

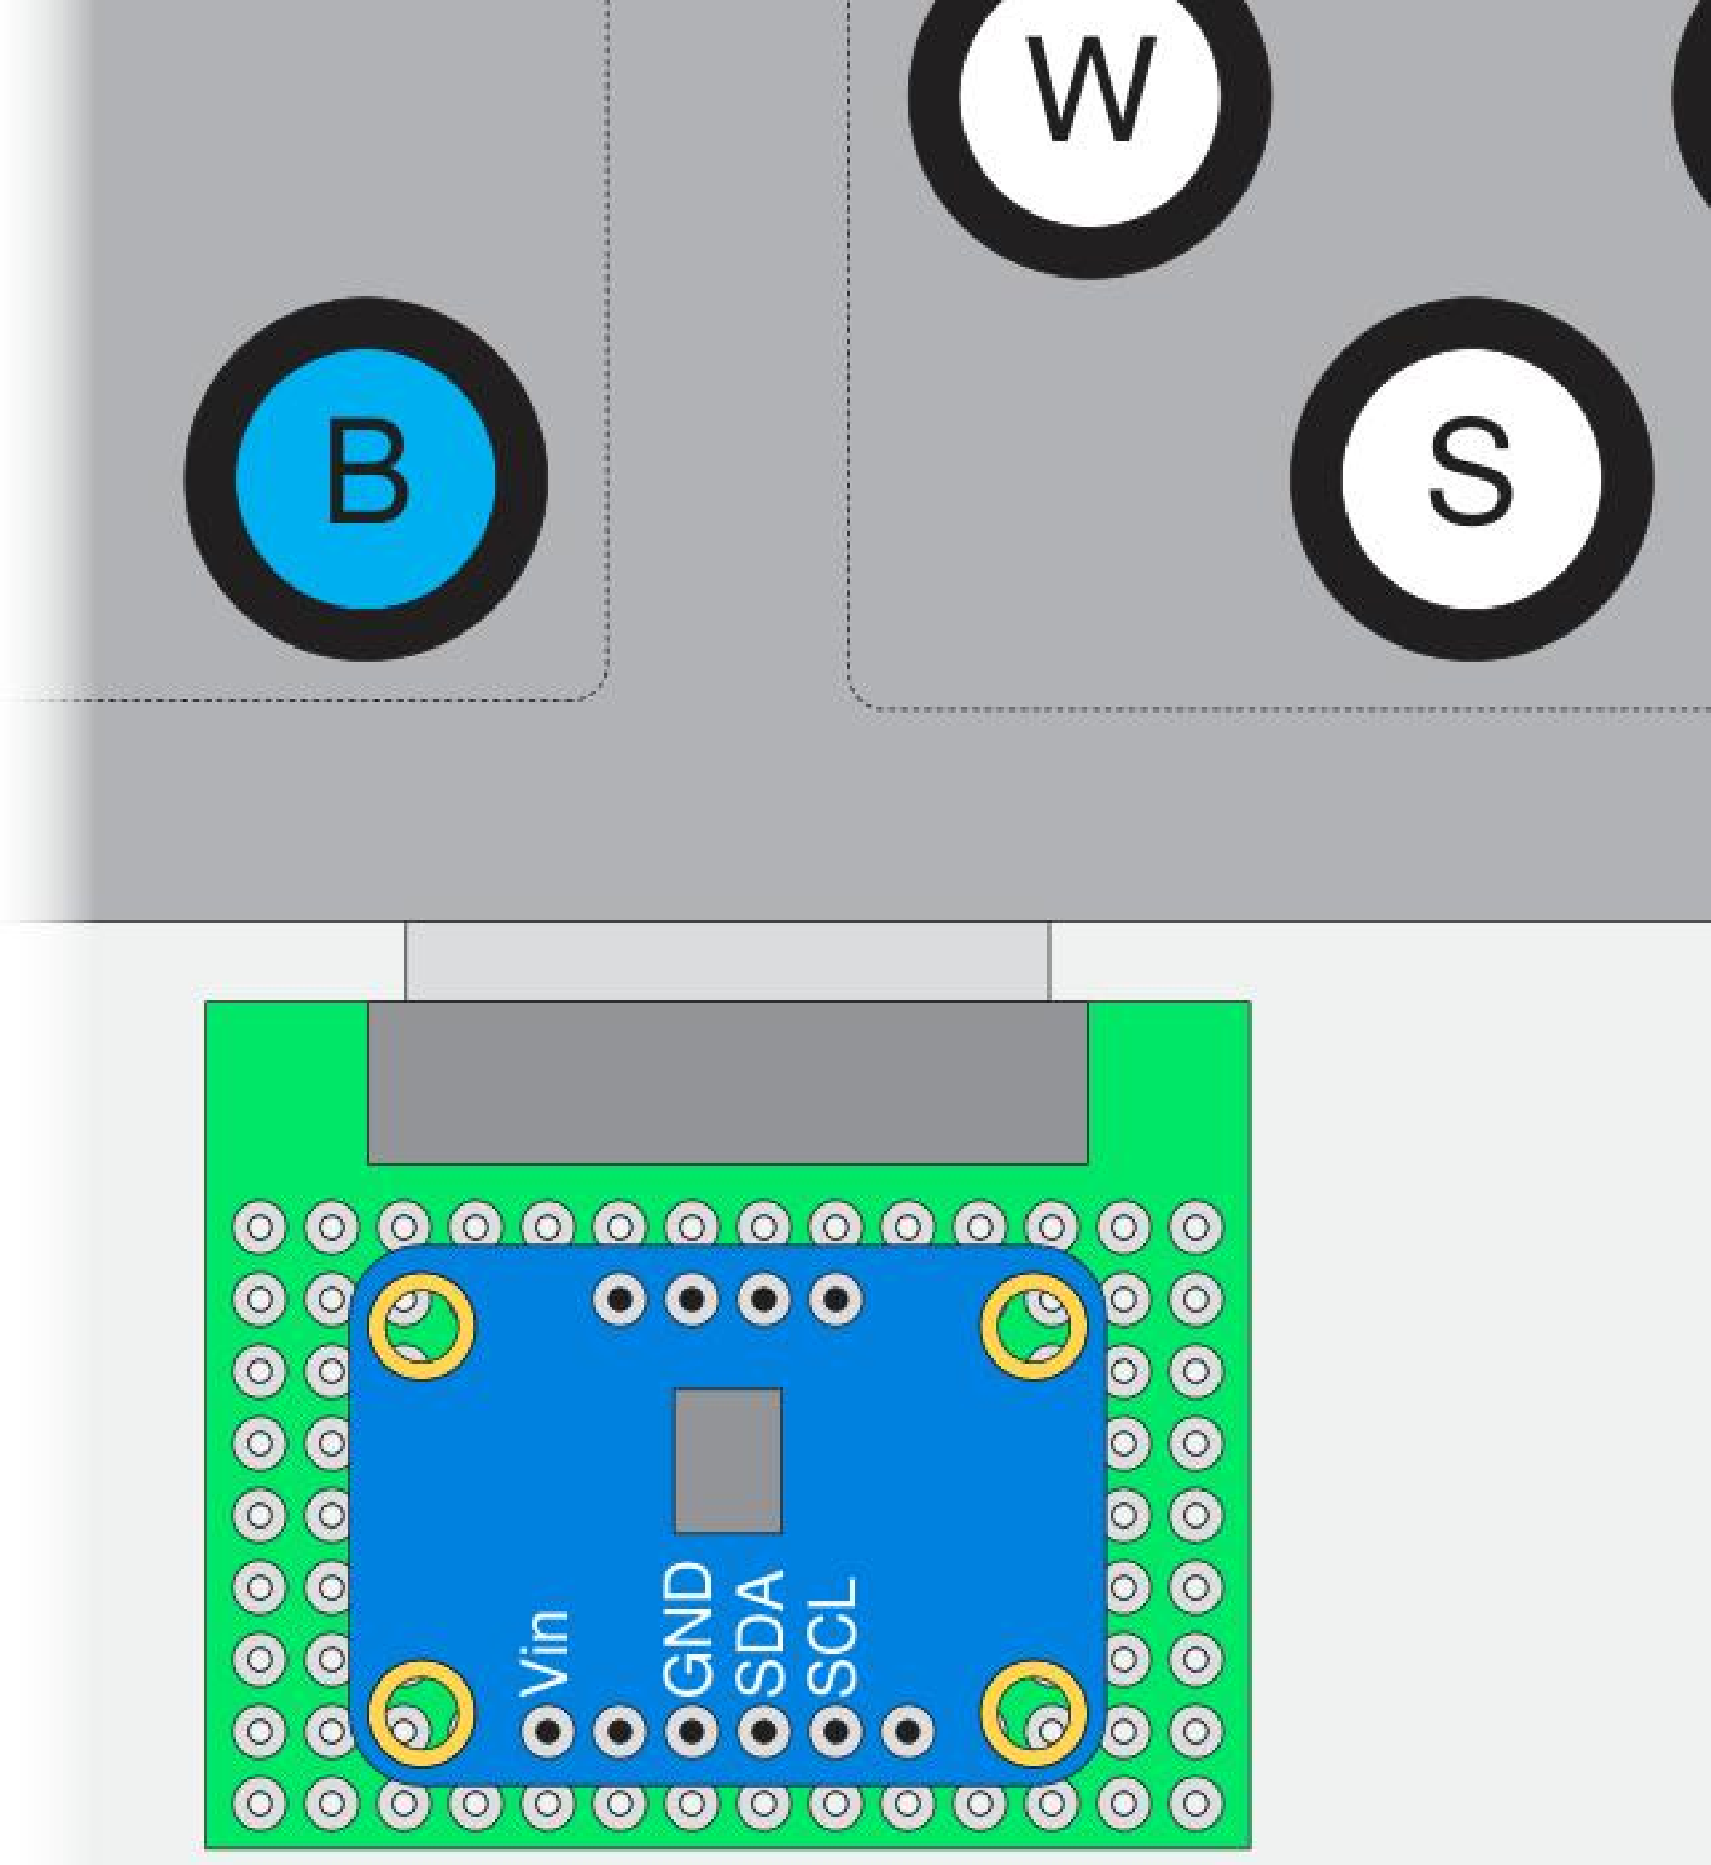

Fig.4: connecting the IMU to the console via its 15-way D-sub connector.

increasingly sensitive and harder to control, which can make the behaviour feel a bit 'twitchy'.

By limiting the range, we keep things well-behaved and easier to use.

Once again, the bulk of our program remains the same (see the file named **CB-jul26-code-03.txt**). One change is that we declare two global variables: a floating-point **G** (the gravitational constant), which we assign the value 9.81, and an integer **MAX_TILT**, which we assign the value 60 (degrees).

Also, we create a simple helper function called **ClampFloat()** that constrains (or 'clamps') a floating-point value to lie within a specified range. This function accepts three floating-point arguments: the value **x**, a lower limit **lo**, and an upper limit **hi**. If **x** is less than **lo**, the function returns **lo**; if **x** is greater than **hi**, it returns **hi**; otherwise, it returns **x** unchanged.

The main changes occur in our **loop()** function, as illustrated in Listing 3(a).

On line 39, we ask the sensor for the current gravity vector, as before.

On lines 42 and 43, we divide the X and Y values by **G** (the strength of gravity, 9.81) to get numbers between about -1 and +1. We then clamp them just in case they stray slightly outside that range due to noise.

On lines 45 and 46, we use the inverse sine (arcsine) function, **asin()**, to convert our **gx** and **gy** values to angles in radians, then multiply by 180/π to convert them into degrees.

On lines 48 and 49, we round the angles to the nearest integer to make them easier to use.

On lines 52 and 53, we limit the values to ±60° to ensure things stay well-behaved and easy to control.

On lines 55 and 56, we print the results.

As usual, as soon as my program was up and running, I tilted my breadboard about 45° to the left, then 45° to the right, then 45° forward, and finally 45° backward. If you examine the results from the full run (in the file named **CB-jul26-results-03.txt**), you'll see that things are now very easy to understand. A drastically cut-down version of these results is shown below:

```
BNO055 tilt angles
LR: left = -ve, right = +ve
FB: back = -ve, front = +ve
LR = -45   FR =  -4   (tilting left)
LR =  45   FR =  -2   (tilting right)
LR =  -1   FR =  45   (tilting forward)
LR =  -1   FR = -45   (tilting backward)
```

Translating to delta values

Having access to specific tilt angles is all well and good, and I'm sure this data will come in handy in the future. For the moment, however, I want to keep things as simple as possible.

Let's say that we are going to use the tilting of our console to 'roll' a virtual marble around our array. Let's also set a tilt threshold of 15°. Anything more than this and the virtual marble will roll; anything less and it will stay where it is.

Since our 'marble' can move by only one pixel at a time, we will make use of integer **deltaX** and **deltaY** values, each of which can be -1, 0, or +1.

As before, the bulk of our program remains unchanged (see the file named **CB-jul26-code-04.txt**). The only modification we make to the program's front end is to add a new global integer variable named **TILT_THRESHOLD**, which we set to 15.

The remaining changes occur in our **loop()** function, as illustrated in Fig. 4(a). On lines 56 and 57, we declare two 8-bit integer variables, **deltaX** and **deltaY**, and set them to 0.

On lines 59-62, we check whether the console is tilted more than 15° to the left or right; if so, we set the **deltaX** value accordingly. Similarly, on lines 64-67, we check whether the console is tilted more than 15° forward or backward; if so, we set the **deltaY** value accordingly.

On lines 69 to 72, we print the actual left-right (LR) tilt angle, followed by

```
56   int8_t deltaX = 0;
57   int8_t deltaY = 0;
58
59   if (leftRightDeg <= -TILT_THRESHOLD)
60     deltaX = -1;
61   else if (leftRightDeg >= TILT_THRESHOLD)
62     deltaX = +1;
63
64   if (frontBackDeg <= -TILT_THRESHOLD)
65     deltaY = -1;
66   else if (frontBackDeg >= TILT_THRESHOLD)
67     deltaY = +1;
68
69   Serial.print("LR = ");   Serial.print   (leftRightDeg);
70   Serial.print("  dX = "); Serial.print   (deltaX);
71   Serial.print("  FB = "); Serial.print   (frontBackDeg);
72   Serial.print("  dY = "); Serial.println (deltaY);
73
74   delay(100);
75 }
```

Listing 4(a): adding tilt threshold and delta values to our game.

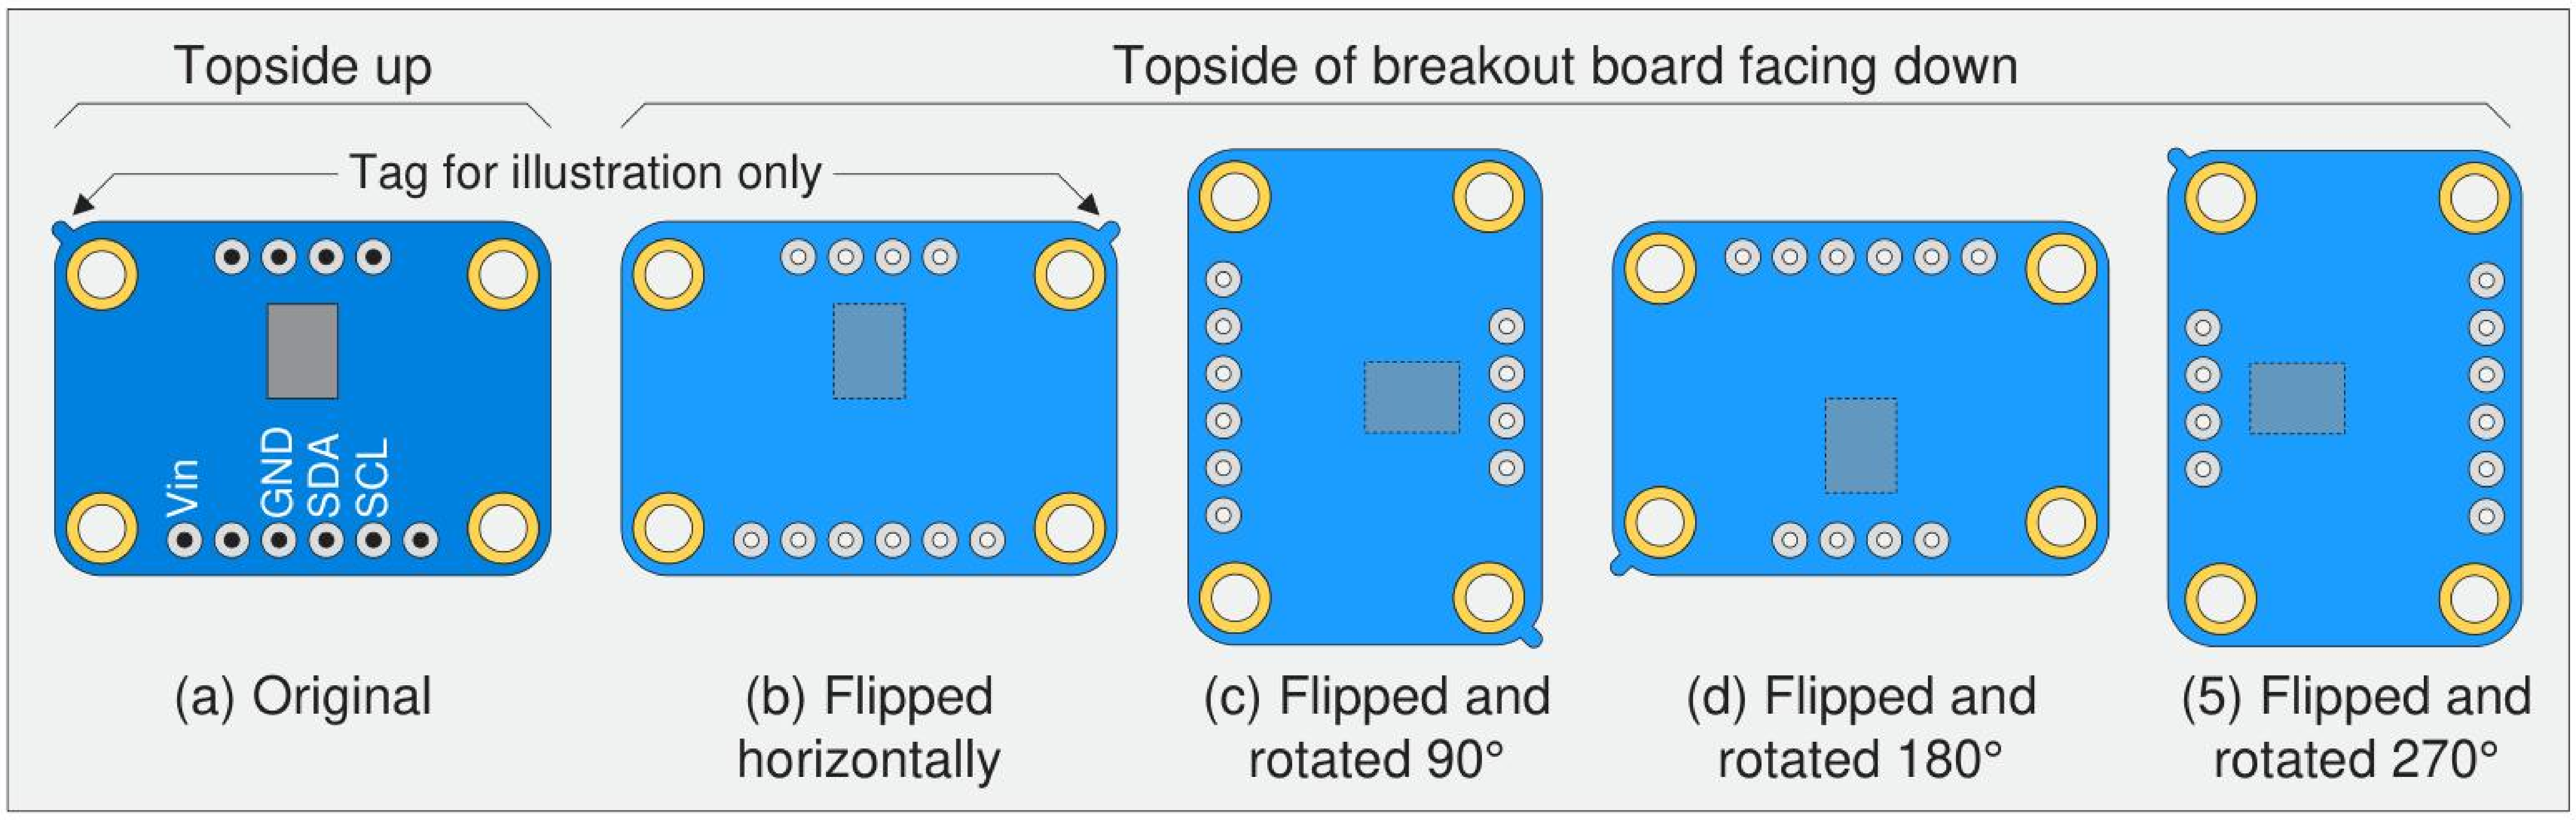
Topside up
Topside of breakout board facing down
Tag for illustration only
Vin
GND
SDA
SCL

(a) Original (b) Flipped horizontally (c) Flipped and rotated 90° (d) Flipped and rotated 180° (5) Flipped and rotated 270°

Fig.5: reflections and rotations depending on our IMU's mounting arrangement.

the corresponding **deltaX** value, then the actual forward-backward (FB) tilt angle, followed by the corresponding **deltaY** value.

Some sample (drastically cut down) results are shown below:

```
BNO055 tilt angles
LR: left = -ve, right = +ve
FB: back = -ve, front = +ve
LR =   0  dX =  0  FB =   0  dY =  0
LR = -18  dX = -1  FB =  -1  dY =  0
LR =  21  dX =  1  FB =   0  dY =  0
LR =   9  dX =  0  FB =  18  dY =  1
LR =   2  dX =  0  FB = -17  dY = -1
```

In this case, the board started out flat, then I tilted it more than 15° left, then right, then forward, then back. Run this yourself and verify that the results are as you expect.

Adding the IMU to the console

This is the exciting part where we get to add the IMU to the console. In the fullness of time, I plan to mount my IMU inside my console. In the short term, however, I've placed it on a small circuit board and connected it via the 15-way D-sub connector, as shown in Fig.4.

Before you do anything else, rerun our latest and greatest "Delta Values" program and verify that everything still works as expected. It will be best to detect and resolve any problems now, before we start adding more bells and whistles.

The main thing to note here is that the IMU is oriented exactly as it was during my earlier experiments. When I eventually mount it inside the console, however, I may flip it upside-down. It could also be rotated by 90°, 180°, or 270°, as illustrated in Fig.5.

The configurations shown in Fig.5(b) and Fig.5(d) will probably require us to invert (change the sign of) one or both of the X and Y values reported by the IMU. The options in Fig.5(c) and Fig.5(e) will also require us to swap the X and Y axes between the IMU and the console (map X to Y and Y to X), in addition to applying any necessary sign changes.

As a rule of thumb, flipping the IMU reverses the direction (sign change), while rotating the IMU reassigns the directions (axis swap).

Rolling a virtual marble

Do you recall the flashing pixel program we created in the March 2026 issue? I'm referring to the one that started with a green pixel flashing in the middle of our array. Pressing any of the red, yellow, green, or blue action buttons changed the pixel's colour, while the north, south, east, and west direction buttons moved it one pixel in that direction.

If the pixel reached the edge of the display, continuing in the same direction caused it to 'wrap around' to the opposite side. You may wish to refresh your memory before we proceed by perusing the program (in the file named **CB-jul26-code-05.txt**) and rerunning it on your console.

How about we keep this functionality while adding the ability to treat the flashing pixel as a virtual marble, tilting the console to 'roll it' around the display? Happily, the way we've designed things means that doing this is going to be quite easy.

We start by copying the four "include library" statements from our latest and greatest 'delta values program' (DVP) (lines 1-4) into our new 'flashing pixel program' (FPP) (lines 5-8 in the file named **CB-jul26-code-06.txt**).

Next, we copy the IMU-related definitions from DVP (lines 6-10) into FPP (lines 10-16). We also copy the **ClampFloat()** function from DVP (lines 12-17) into the 'game generic routines' section of FPP (lines 439-445).

In the 'game related stuff' section of FPP, just under the 1ms **TICK** definition (line 151), we add a 200ms **IMU_SAMPLE_TICKS** definition (line 152). We also add a new **MarbleThen** variable (line 182). We'll use these to periodically sample the IMU and update the marble's position.

To initialise the IMU in the **setup()** function, we copy the key statements from DVP (lines 29, 34, and 35) into FPP (lines 213-216). In the original FPP **loop()** function, we repeatedly called our **UpdateFlash()**, **UpdateSwitches()** and **ProcessSwitches()** functions each millisecond. The only change we're making here is to follow these with a call to a new **UpdateMarbleFromIMU()** function.

Essentially, we are going to take the entire **loop()** function from our DVP, chop out anything we don't need, augment it with just a *soupçon* of additional logic, and add it into FPP as the aforementioned **UpdateMarbleFromIMU()** function (lines 499-537). Let's take a moment to touch on some of the highlights, as illustrated in Listing 6(a).

On line 501, we check whether sufficient time has passed since we last sampled the IMU and moved the marble (this delay, as specified by **IMU_SAMPLE_TICKS**, is currently set to 200ms, but you can increase or decrease it as you wish). If not, we simply skip the rest of this function and try again on the next iteration of the main loop.

The bulk of this function (lines 503-530) is lifted from the DVP, ultimately yielding **deltaX** and **deltaY** values of -1, 0, or +1.

On lines 532–533, we perform a quick sanity check: if both **deltaX** and **deltaY** are zero, the board is essentially level (or close enough), and there's nothing to do. However, if either value is non-zero—meaning we've tilted beyond our threshold in at least one direction—we call our trusty **UpdateLocation()** function to move the marble by one step in the appropriate direction.

Finally, on line 535, we store the current time in **MarbleThen**, thereby marking the time of this marble update. The next update won't occur until another IMU_SAMPLE_TICKS milliseconds have elapsed.

As you'll see when you run this new version of our program, we can still change the flashing pixel's colour using the action buttons and move it with the direction buttons, as before. However, we can also 'roll the marble' by tilting the console.

You'll observe that, in addition to rolling the marble north, south, east or west, we can move it along the diagonals—northeast, southeast, northwest or southwest—a trick we never bothered to implement under button control.

Returning to the Countdown Timer

Ironically and paradoxically, it seems we've run out of time, so we'll have to regroup on this next month.

As always, if you have any thoughts on anything you've read here, please feel free to drop me an email at max@clivemaxfield.com PE

```
void UpdateMarbleFromIMU()
{
  if ( (TimeNow - MarbleThen) >= IMU_SAMPLE_TICKS)
  {
    imu::Vector<3> grav = IMU.getVector(Adafruit_BNO055::VECTOR_GRAVITY);

    // Normalize and protect against tiny over/under-shoots
    float gx = ClampFloat(grav.x() / G, -1.0, 1.0);
    float gy = ClampFloat(grav.y() / G, -1.0, 1.0);

    float leftRightDegF = asin(gx) * 180.0 / PI;
    float frontBackDegF = asin(gy) * 180.0 / PI;

    int leftRightDeg = (int)round(leftRightDegF);
    int frontBackDeg = (int)round(frontBackDegF);

    // Optional clipping for display / gameplay
    leftRightDeg = constrain(leftRightDeg, -MAX_TILT, MAX_TILT);
    frontBackDeg = constrain(frontBackDeg, -MAX_TILT, MAX_TILT);

    int8_t deltaX = 0;
    int8_t deltaY = 0;

    if (leftRightDeg <= -TILT_THRESHOLD)
      deltaX = -1;
    else if (leftRightDeg >= TILT_THRESHOLD)
      deltaX = +1;

    if (frontBackDeg <= -TILT_THRESHOLD)
      deltaY = -1;
    else if (frontBackDeg >= TILT_THRESHOLD)
      deltaY = +1;

    if ((deltaX != 0) || (deltaY != 0))
      UpdateLocation(deltaX, deltaY);

    MarbleThen = millis();
  }
}
```

Listing 6(a): our new virtual marble function.

Devices and components used in this issue

9DOF Adafruit IMU	www.amazon.co.uk/dp/B017PEIGIG

ETI BUNDLE (1) Teach-In 3, 4 and 5 – all on CD-ROM – only £18.95

ELECTRONICS TEACH-IN 3 – CD-ROM
Mike & Richard Tooley

The three sections of the *Teach-In 3* CD-ROM cover a huge range of subjects that will interest everyone involved in electronics – from newcomers to the hobby and students to experienced constructors and professionals.

The first section (80 pages) is dedicated to *Circuit Surgery*, *EPE/PE's* regular clinic dealing with readers' queries on circuit design problems – from voltage regulation to using SPICE circuit simulation software.

The second section – *Practically Speaking* – covers hands-on aspects of electronics construction. Again, a whole range of subjects, from soldering to avoiding problems with static electricity and identifying components is covered. Finally, our collection of *Ingenuity Unlimited* circuits provides over 40 circuit designs submitted by readers.

The CD-ROM also contains the complete *Electronics Teach-In 1* book, which provides a broad-based introduction to electronics in PDF form, plus interactive quizzes to test your knowledge and TINA circuit simulation software (a limited version – plus a specially written TINA Tutorial). *Teach-In 1* covers everything from electric current through to microprocessors and microcontrollers, and each part includes demonstration circuits to build on breadboards or to simulate on your PC.

ELECTRONICS TEACH-IN 4 – CD-ROM, A BROAD INTRODUCTION TO ELECTRONICS
Mike & Richard Tooley

The *Teach-In 4* CD-ROM covers three of the most important electronics units that are currently studied in many schools and colleges. These include, Edexcel BTEC level 2 awards and the electronics units of the Diploma in Engineering, Level 2.

The CD-ROM also contains the full *Modern Electronics Manual*, worth £29.95. **The Manual** contains over 800 pages of electronics theory, projects, data, assembly instructions and web links. It's an ideal way to learn about electronics.

ELECTRONICS TEACH-IN 5 – CD-ROM
JUMP START
Mike & Richard Tooley

15 design and build circuit projects for newcomers or those following courses in school and colleges.

The projects are: ■ Moisture Detector ■ Quiz Machine ■ Battery Voltage Checker ■ Solar-Powered Charger ■ Versatile Theft Alarm ■ Spooky Circuits ■ Frost Alarm ■ Mini Christmas Lights ■ iPod Speaker ■ Logic Probe ■ DC Motor Controller ■ Egg Timer ■ Signal Injector Probe ■ Simple Radio Receiver ■ Temperature Alarm.

PLUS

PIC'n' Mix – starting out with the popular range of PIC microcontrollers and *Practically Speaking* – tips and techniques for project construction.

The CD-ROM also contains:

- ■ Complete *Teach-In 2* book, a practical introduction to PIC microprocessors
- ■ MikroElektronika, Microchip and L-Tek PoScope software.

Three Teach-ins for the great price of £18.95

PLUS you also get the contents of the free CD-ROM from each issue...

... so that's another TWO Teach-Ins and The Full Modern Electronics Manual!

Teach-In 2026

by Mike Tooley

World of Wireless – An Introduction to Radio and Wireless Technology, part 12.9

RF sources and oscillators

Last month, our focus shifted to frequency conversion methods used in both receivers and transmitters, with particular emphasis on how multipliers and mixers provide us with a means of changing the frequency of a signal.

This month, we take a deep dive into RF sources, explaining how to get the best out of a variety of phase-locked loop (PLL) signal and marker/generator modules, plus direct digital synthesis (DDS) chips. Our *Hands-On* constructional project is an accurate calibration source and programmable frequency marker.

Free-running RF oscillators

In the December 2025 issue, the part 2 article ("Introducing RF components and circuits") briefly mentioned some basic RF oscillators that use quartz crystals as the frequency-determining element. These circuits included the popular Pierce, Colpitts, Miller and Franklin oscillator configurations.

Since there's no attempt to regulate the output frequency, these arrangements can all be described as 'free running'. Free-running oscillators are found in many applications, but their use is constrained by the quality and accuracy of the frequency-determining elements, including the effects of temperature, supply voltage and ageing.

Before we move on to investigate oscillator stability and methods of frequency control, let's look at one more type of free-running oscillator that offers several advantages over the Colpitts oscillator that we've previously used.

The Franklin oscillator

Earlier in the series, we used a simple Colpitts oscillator to generate RF signals over a relatively narrow tuning range. For applications needing a much wider tuning range and good frequency stability, you might prefer the Franklin oscillator configuration (see Fig.9.1).

It requires two separate stages, making it a little more elaborate than the Colpitts arrangement. The advantage is that the tuned circuit components (L1 and C1) can be very loosely coupled to the two active devices.

The circuit employs two N-channel junction gate field-effect transistors (JFETs), TR1 and TR2, with a third insulated-gate FET (Mosfet TR3) configured as a source follower to improve isolation and supply enough current to drive a relatively low-impedance load at the output.

The Franklin oscillator employs a conventional parallel resonant circuit composed of L1 and C1. To improve frequency stability, the resonant circuit is minimally coupled to the active components through low-value capacitors C2 and C3. With appropriate changes to component values, the arrangement shown in Fig.9.1 is suitable for operation from below 100kHz to well over 50MHz.

Fig.9.2 shows a practical realisation of Fig.9.1. Where C1 is a conventional variable capacitor that can

Fig.9.1: the wide-range Franklin oscillator arrangement.

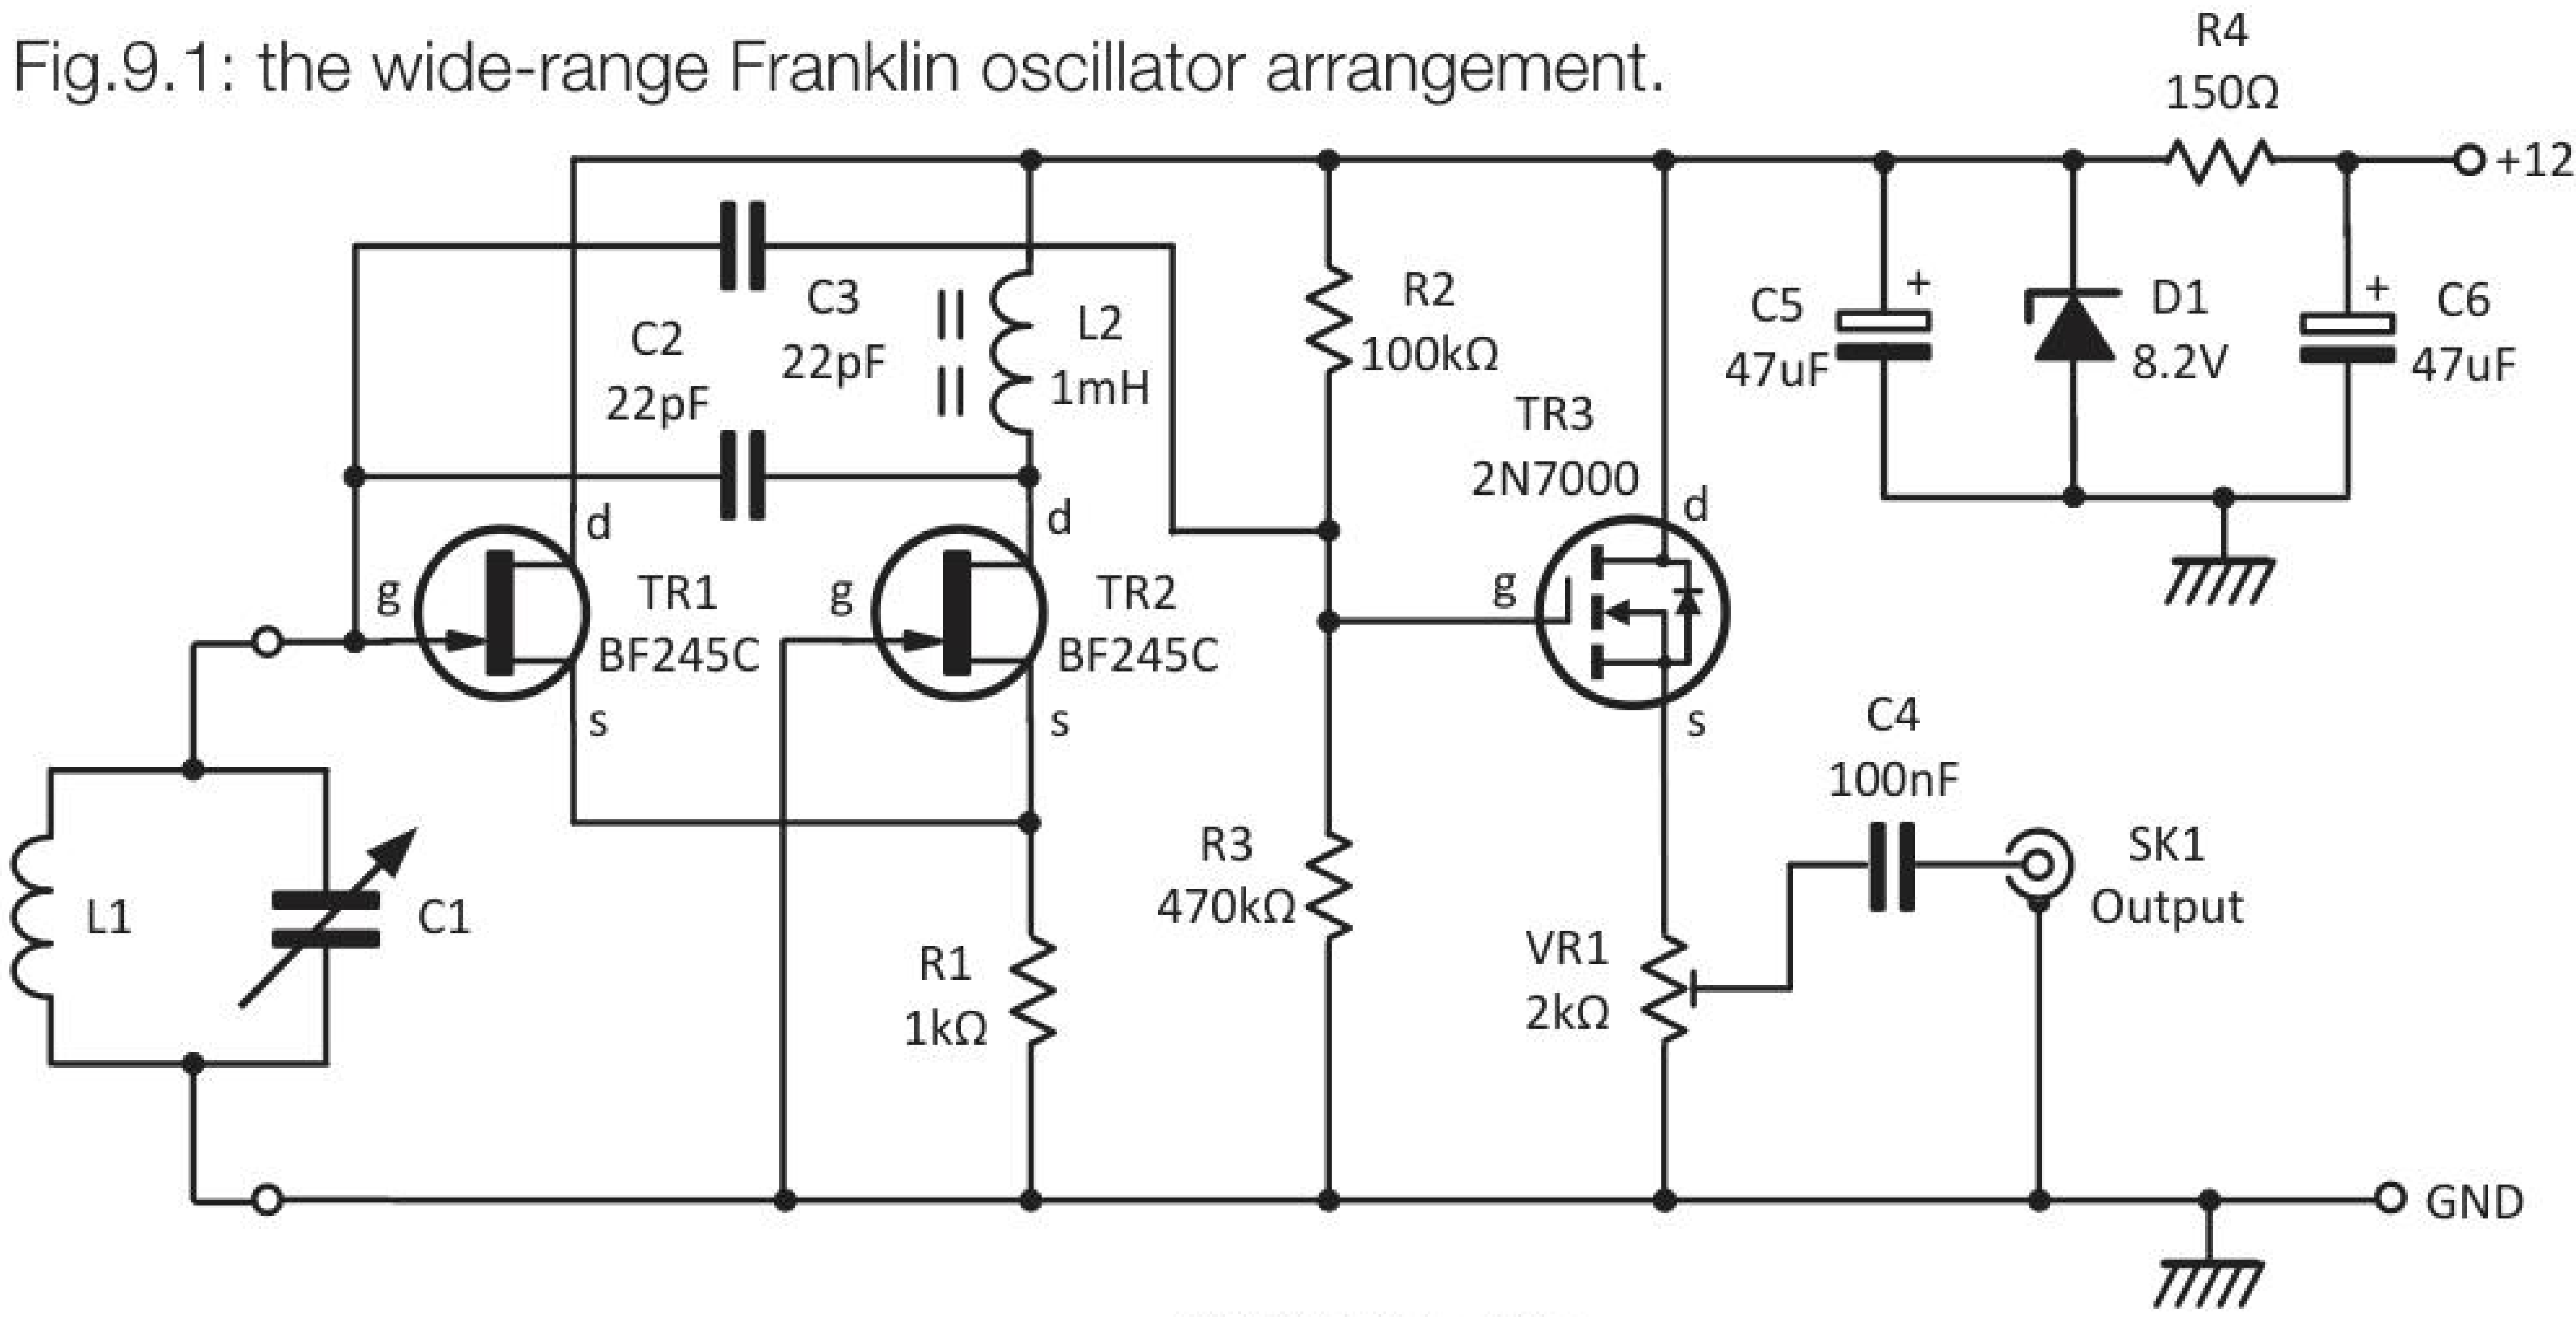

Fig.9.2: a practical realisation of Fig.9.2 (L1 and C1 are external).

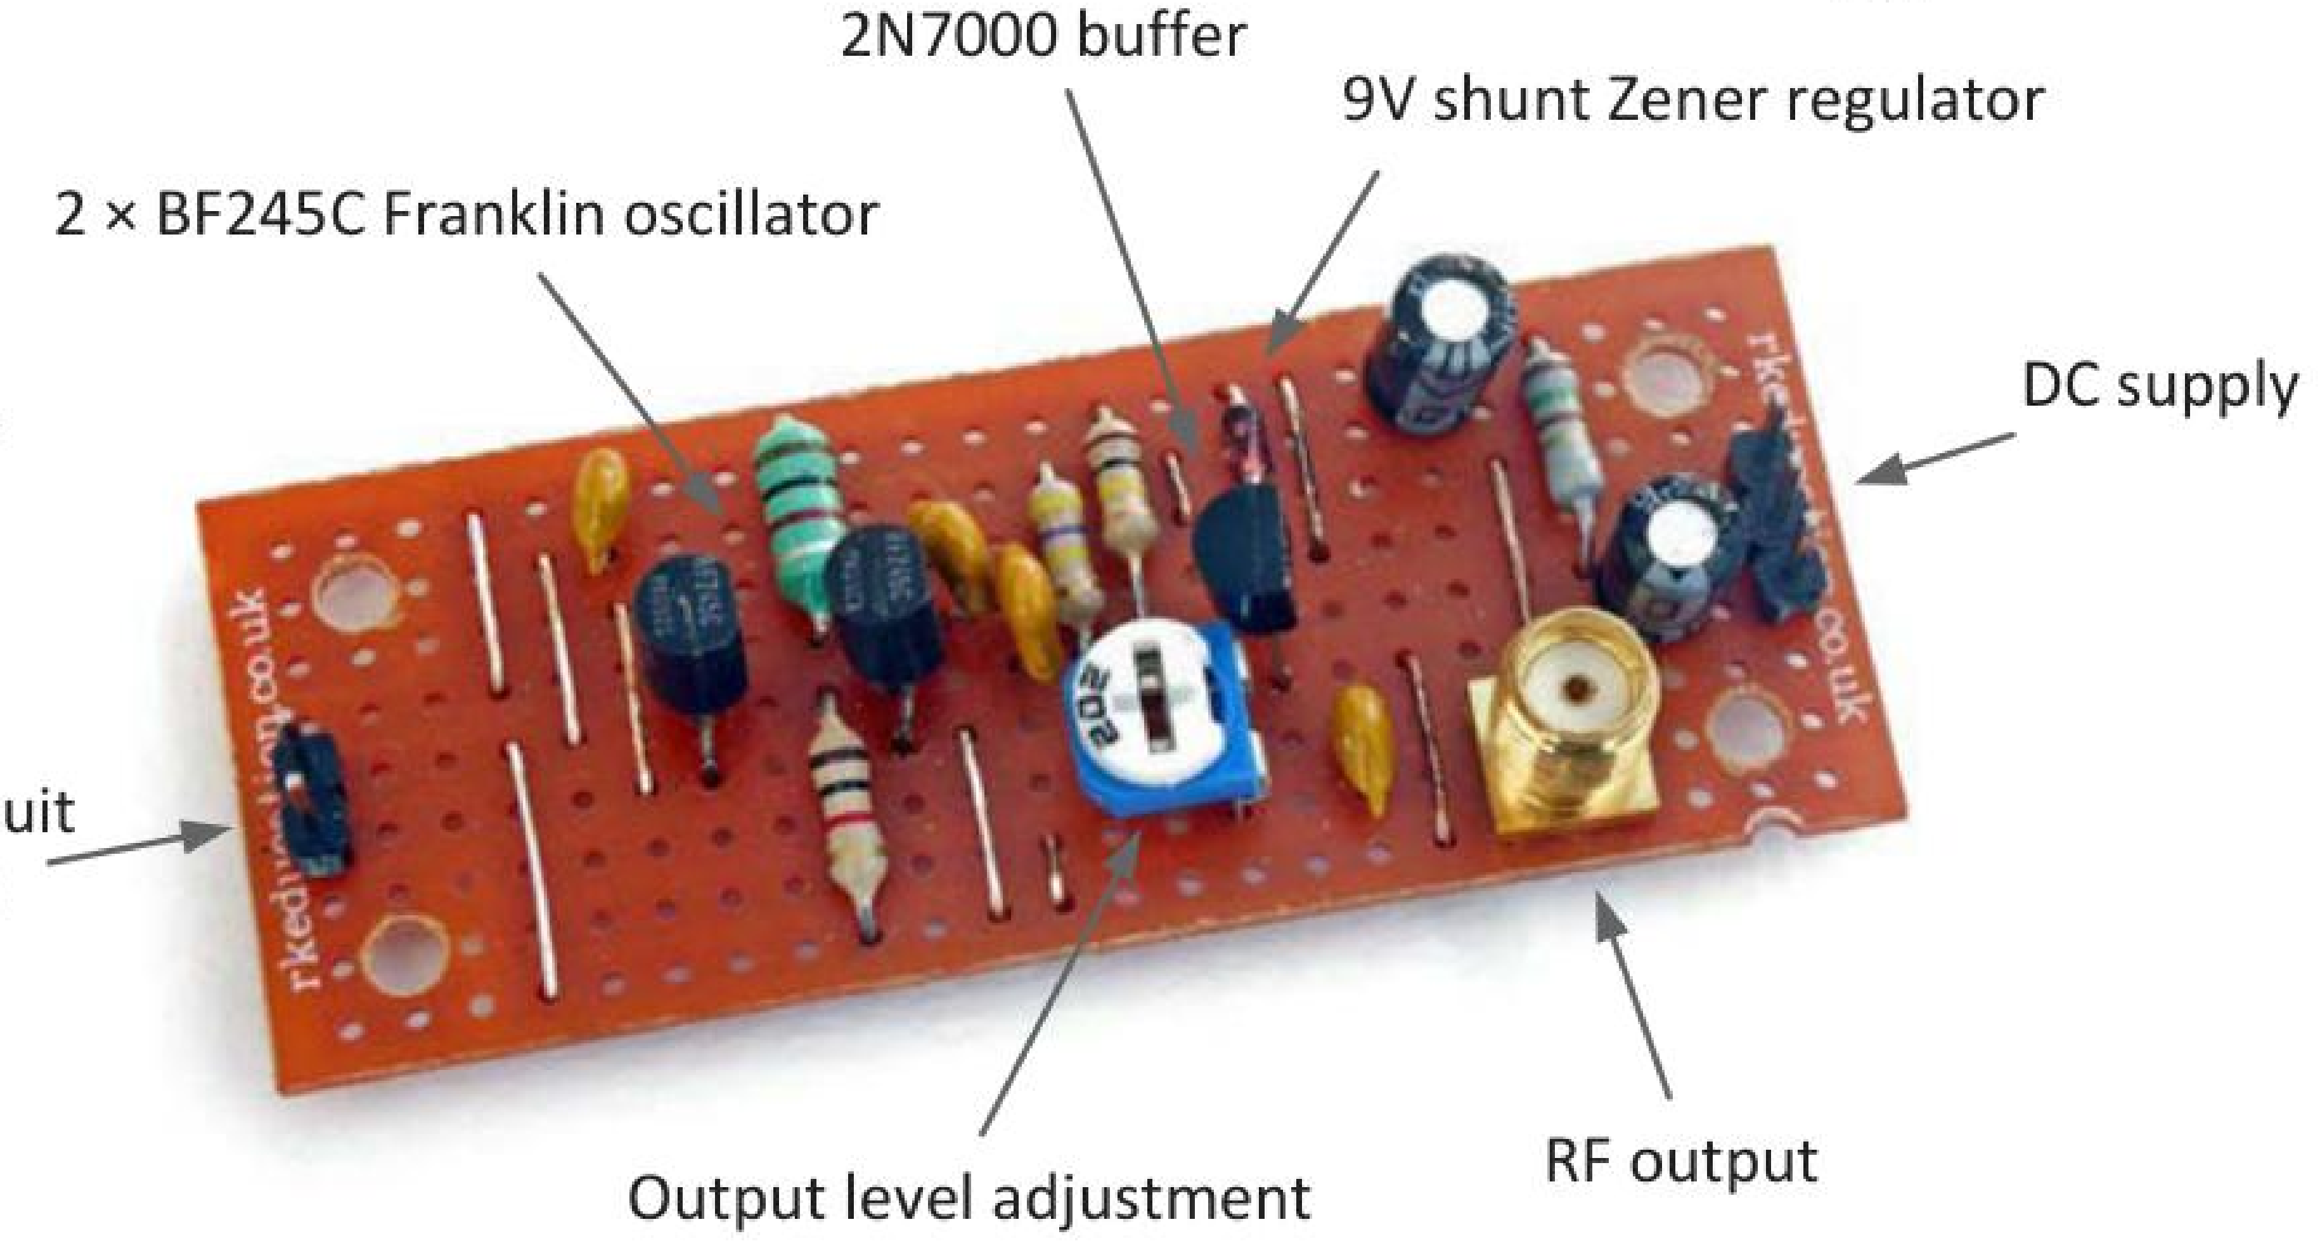

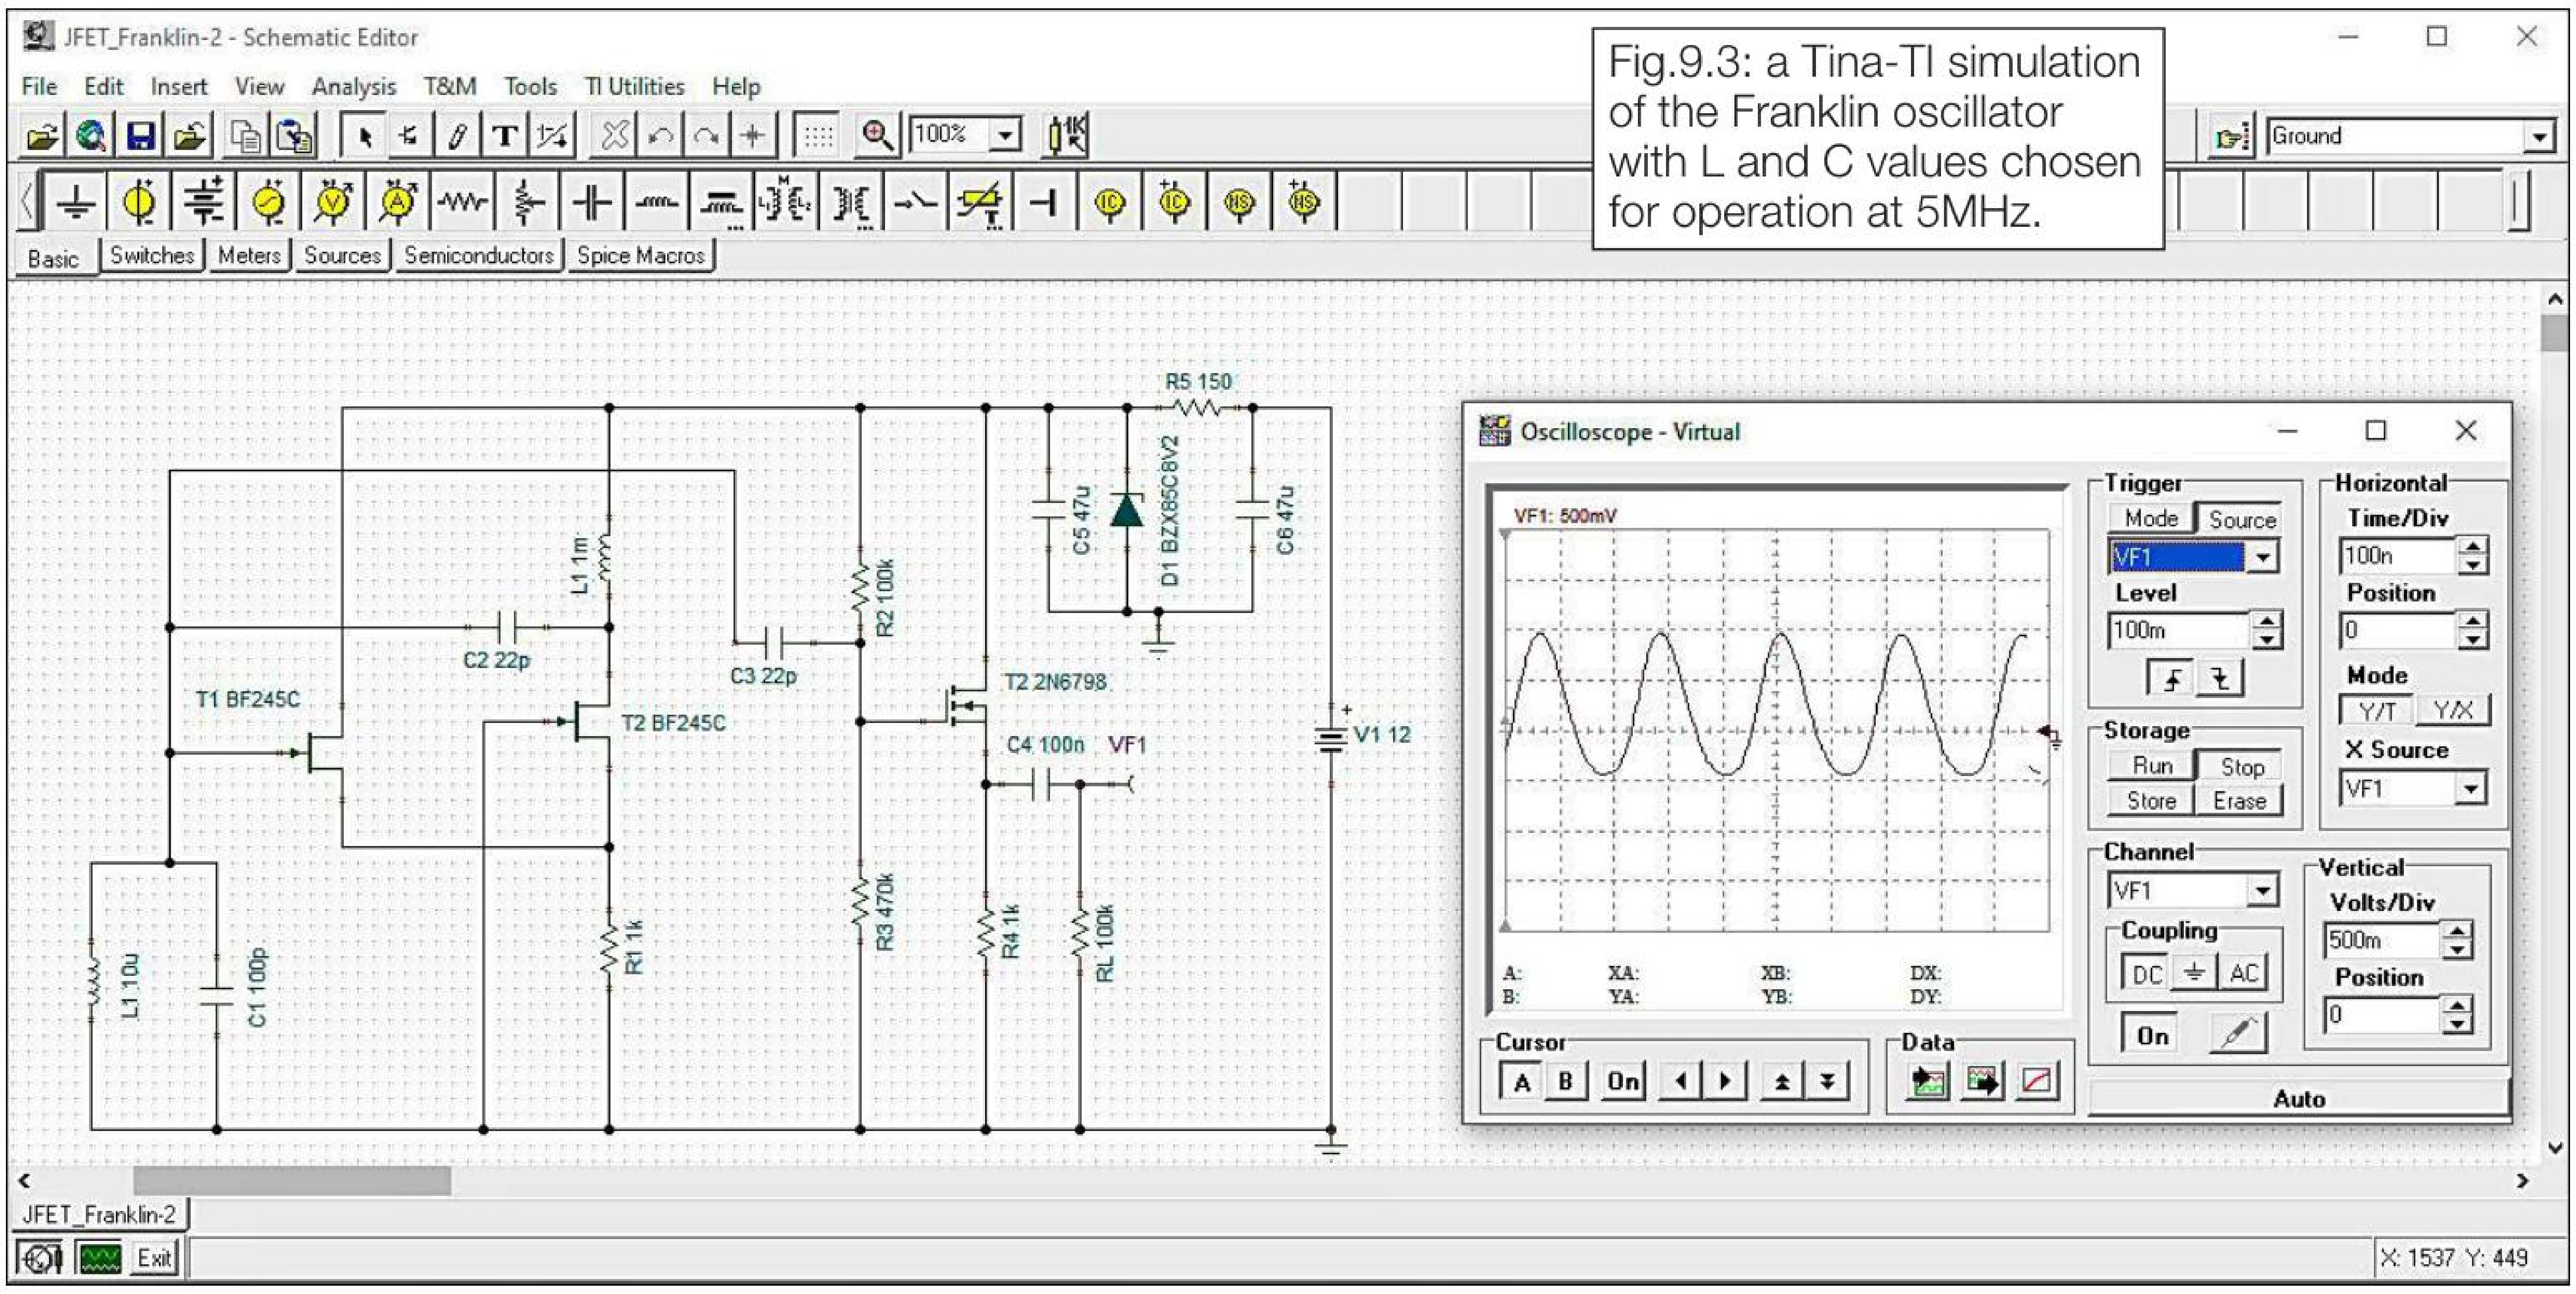
JFET_Franklin-2 - Schematic Editor
File Edit Insert View Analysis T&M Tools TI Utilities Help
Basic Switches Meters Sources Semiconductors Spice Macros
Ground
R5 150
C5 47u
D1 BZX85C8V2
C6 47u
L1 1m
C2 22p
C3 22p
R2 100k
T1 BF245C
T2 BF245C
T2 2N6798
C4 100n
VF1
V1 12
L1 10u
C1 100p
R1 1k
R3 470k
R4 1k
RL 100k
Oscilloscope - Virtual
VF1: 500mV
Trigger
Mode
Source
Level
Storage
Run
Stop
Store
Erase
Channel
Coupling
DC
AC
On
Horizontal
Time/Div
Position
Mode
Y/T
Y/X
X Source
Vertical
Volts/Div
Position
Cursor
A B On
Data
Auto
A: B: XA: YA: XB: YB: DX: DY:
JFET_Franklin-2
Exit
X: 1537 Y: 449

Trigger Source: VF1
Level: 100m
Time/Div: 100n
Position: 0
X Source: VF1
Channel: VF1
Volts/Div: 500m
Position: 0

Fig.9.3: a Tina-TI simulation of the Franklin oscillator with L and C values chosen for operation at 5MHz.

be adjusted from 10pF to 290pF, the values required for L1 (not shown in Fig.9.2) can be determined from Table 9.1. This makes the circuit ideal for use as the basis of a general-purpose RF signal generator with several switched ranges.

Fig.9.3 shows the Tina-TI simulation of the Franklin oscillator in Fig.9.1. For simplicity, we've replaced the preset output amplitude control (RV1) with a fixed resistor. The values of C1 and L1 (100pF and 10μH respectively) were chosen so that the tuned circuit is resonant at approximately 5MHz.

Tina's virtual oscilloscope indicates that the circuit produces an output of approximately 1.5V peak-to-peak with a period of around 200ns.

Tina-TI can be used to investigate the working frequency range of the Franklin oscillator in Fig.9.3. Start by experimenting with the values of L1 and C1 (see Table 9.1) and also try reducing the values of the coupling capacitors, C2 and C3, to provide greater isolation for the resonant circuit.

It is informative to observe the effect of reducing the load impedance, RL, on the output waveform. When the output voltage has fallen to half its open-circuit value, the oscillator's output impedance will be equal to the value of RL.

To obtain a specific output frequency from the circuit in Fig.9.3, you can determine the values of L1 and C1 using the formula on page 4 of the December 2025 issue: $f_0 = 1 \div 2\pi\sqrt{LC}$. Click on L1 and C1 in the Tina simulation to change their values and observe the effect on the output period and waveform.

Electronic tuning

Up to now, our variable RF oscillators have all relied on mechanical tuning of the variable capacitor used in the frequency-determining LC resonant circuit. Variable capacitors can be large and costly; they're also unsuitable for situations where automatic frequency control (AFC) is needed.

A different approach is to use electronic tuning based on a variable capacitance diode (also known as a 'varicap' or 'varactor').

Figs.9.4 & 9.5 show a Colpitts oscillator with electronic tuning via a variable capacitance diode. The tuned circuit is L1 (a 3.3μH inductor) in parallel with the series combination of C2 (a 10nF capacitor) and D1 (a 1SV149 variable capacitance diode).

Table 9.1: inductances for Franklin oscillator with 10-290pF capacitance

Inductance	Approximate tuning range
150μH	0.9-2.5MHz
33μH	1.8-5.2MHz
6.8μH	3.2-12.4MHz
3.3μH	5.2-16.5MHz
1μH	9.3-27.6MHz

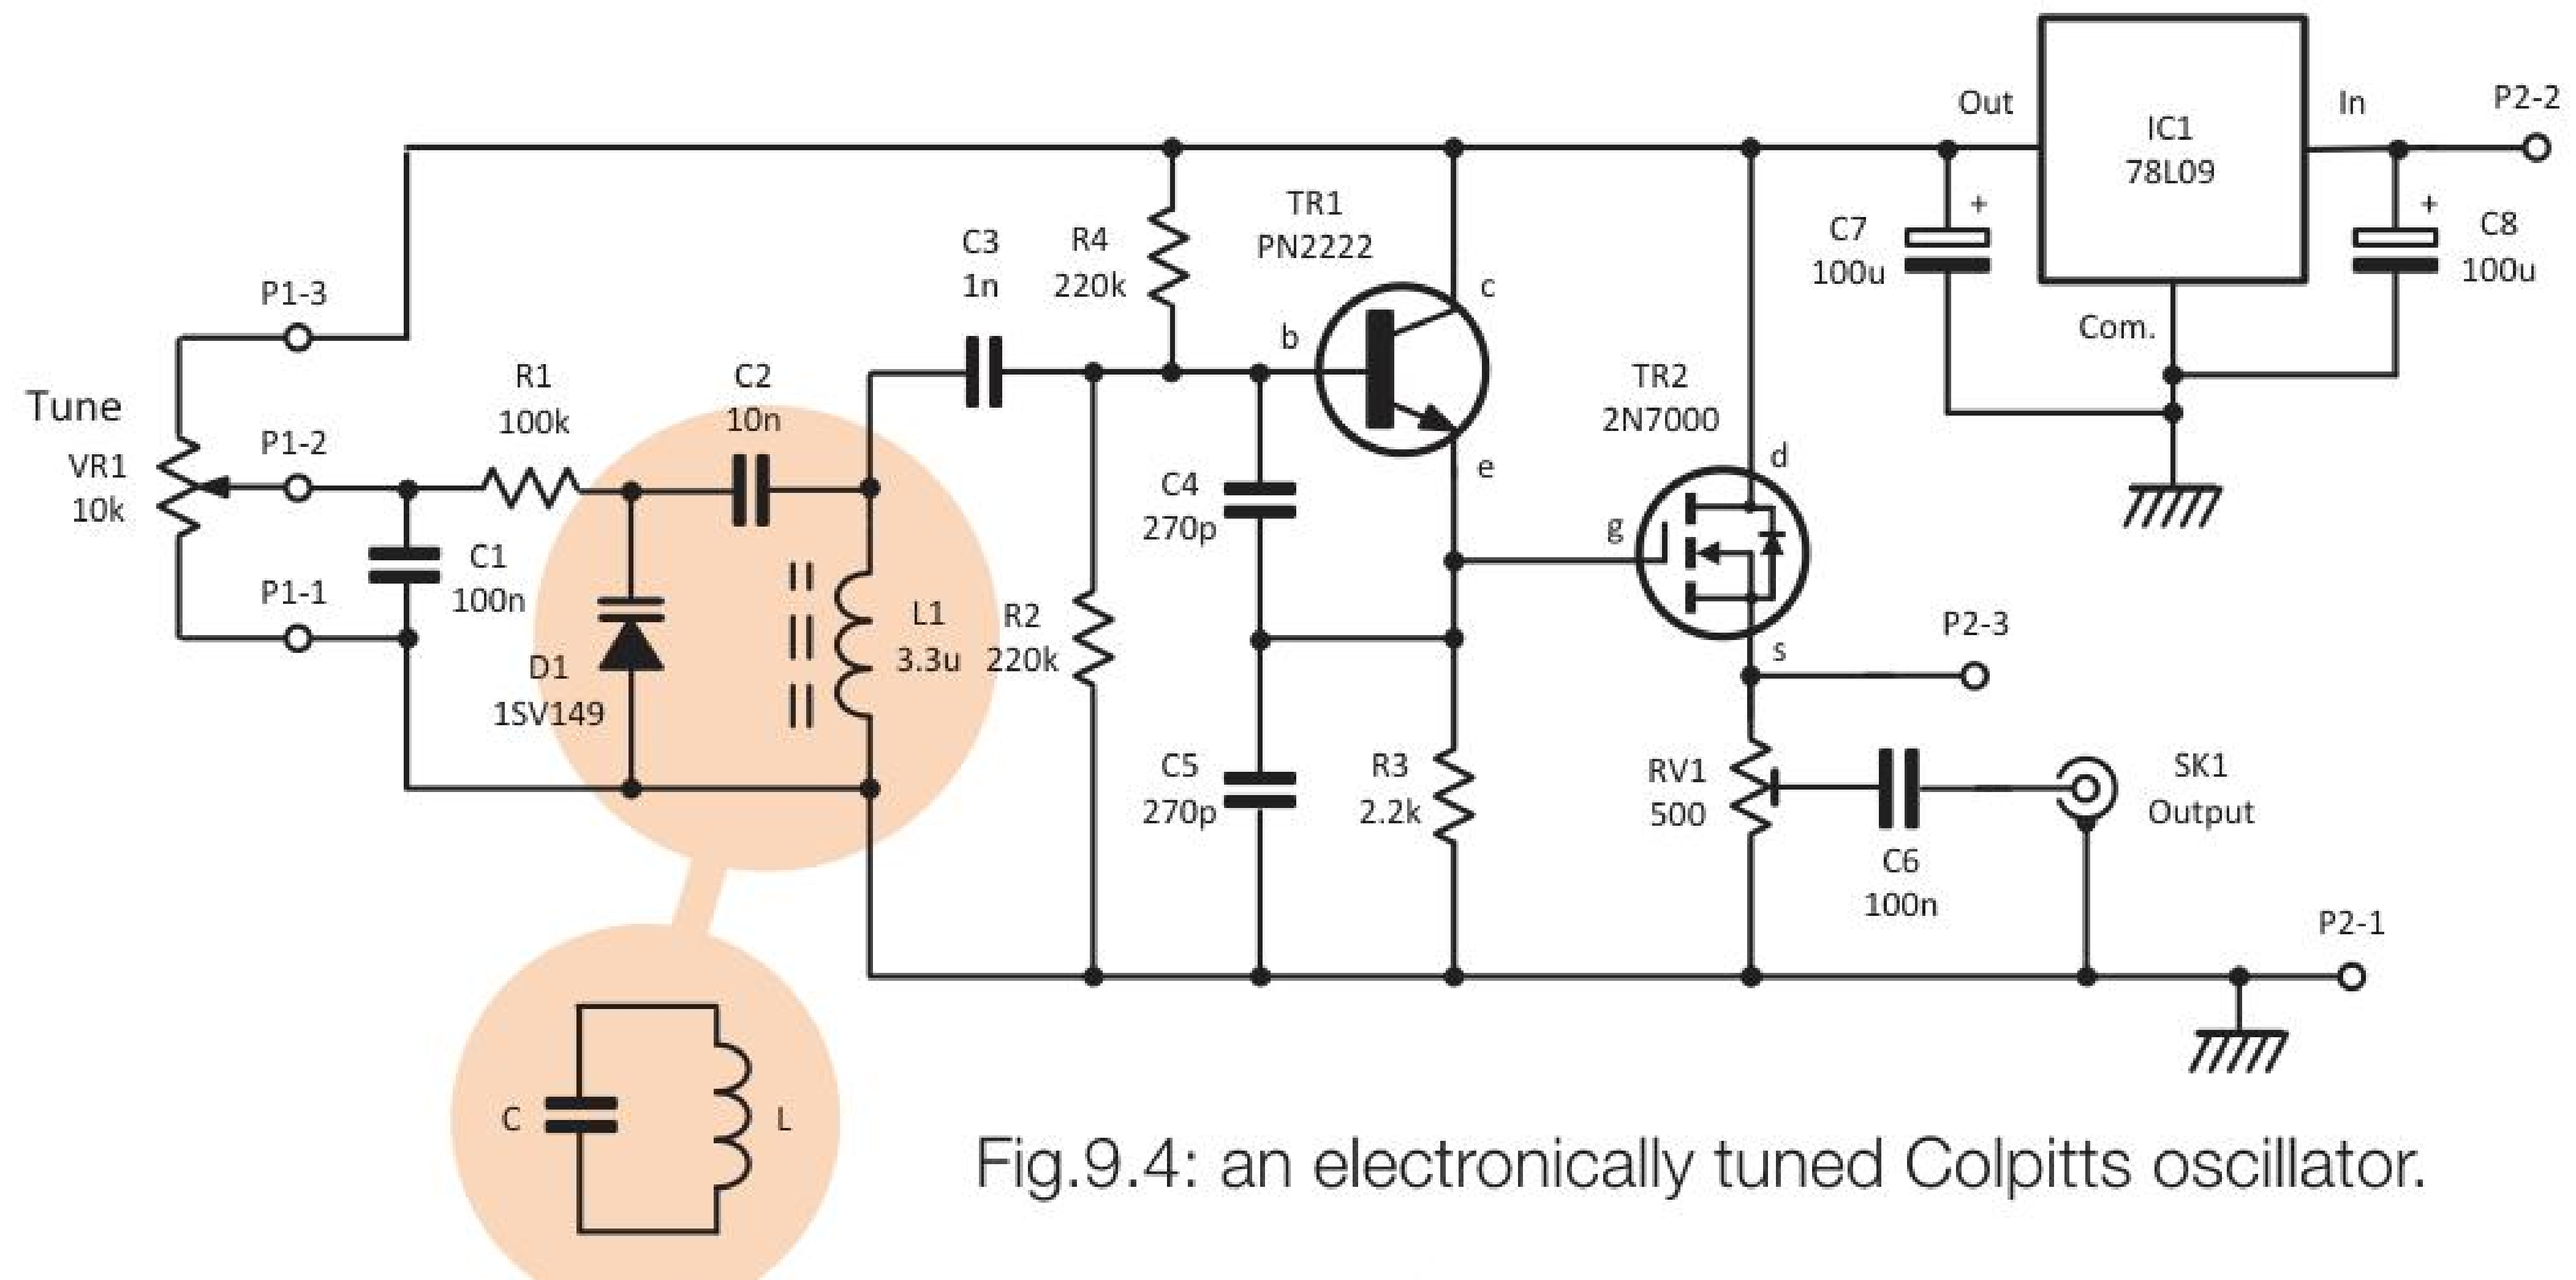
Out
IC1
78L09
In
P2-2
C7
100u
C8
100u
Com.
TR1
PN2222
C3
1n
R4
220k
b
c
e
P1-3
Tune
VR1
10k
P1-2
R1
100k
C2
10n
P1-1
C1
100n
D1
1SV149
L1
3.3u
R2
220k
C4
270p
C5
270p
R3
2.2k
TR2
2N7000
d
g
s
P2-3
RV1
500
C6
100n
SK1
Output
P2-1
C
L

Fig.9.4: an electronically tuned Colpitts oscillator.

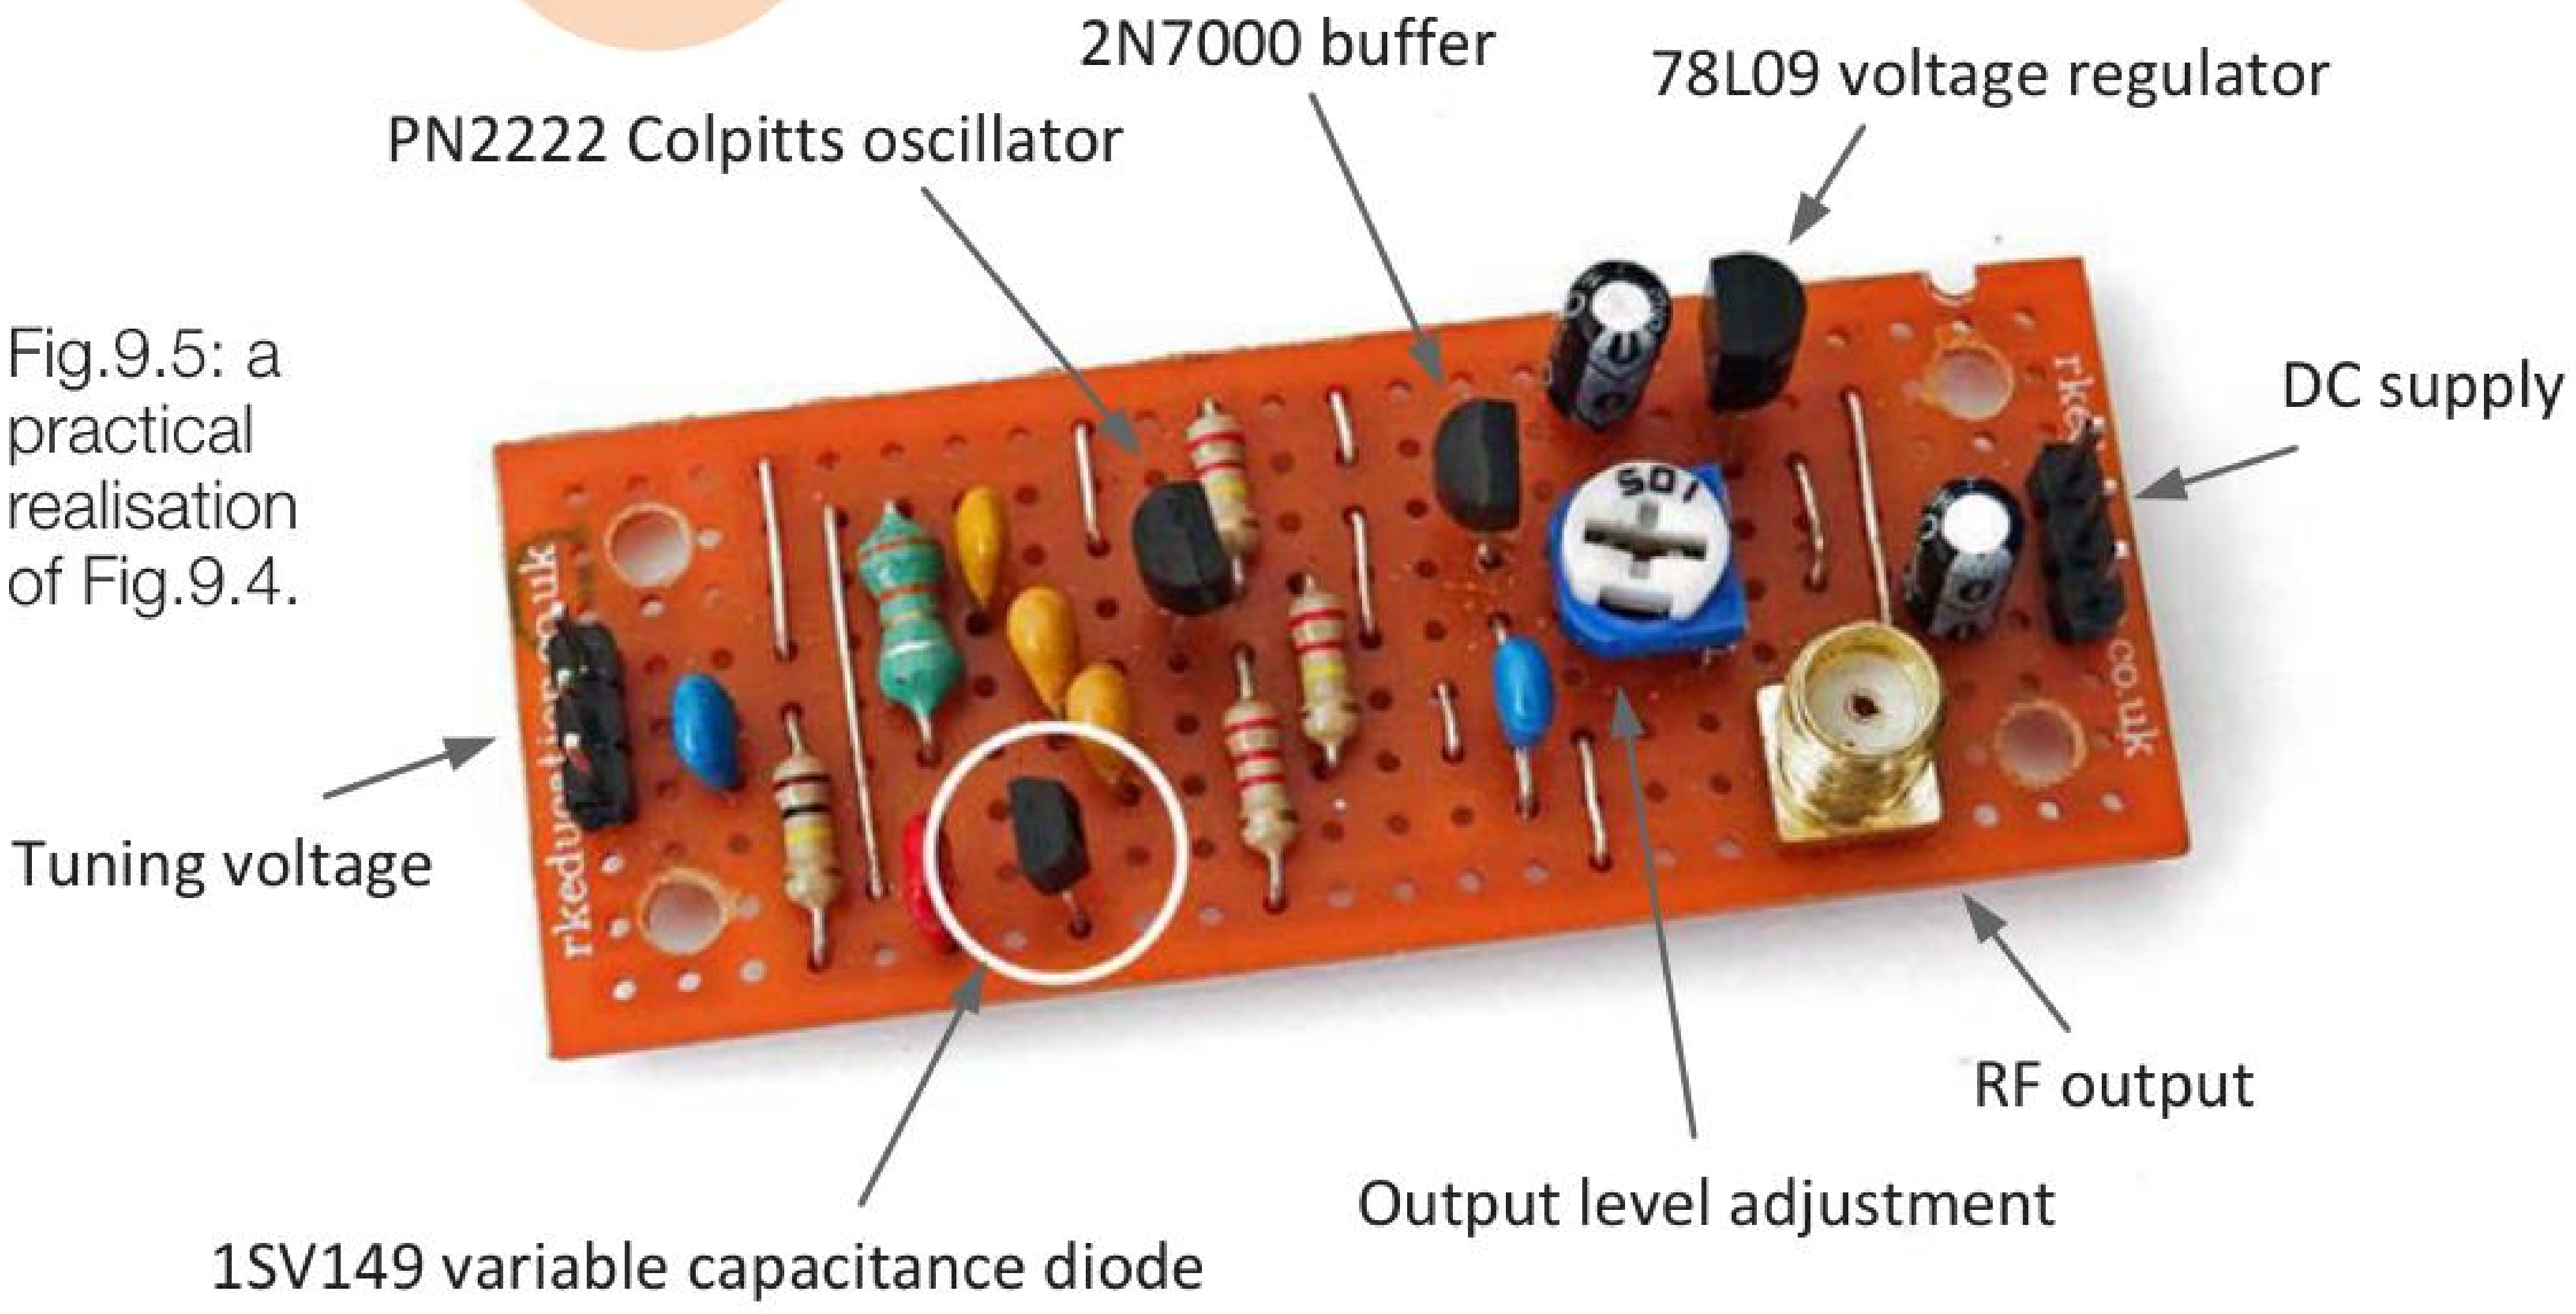
2N7000 buffer
PN2222 Colpitts oscillator
78L09 voltage regulator
DC supply
Tuning voltage
RF output
Output level adjustment
1SV149 variable capacitance diode

Fig.9.5: a practical realisation of Fig.9.4.

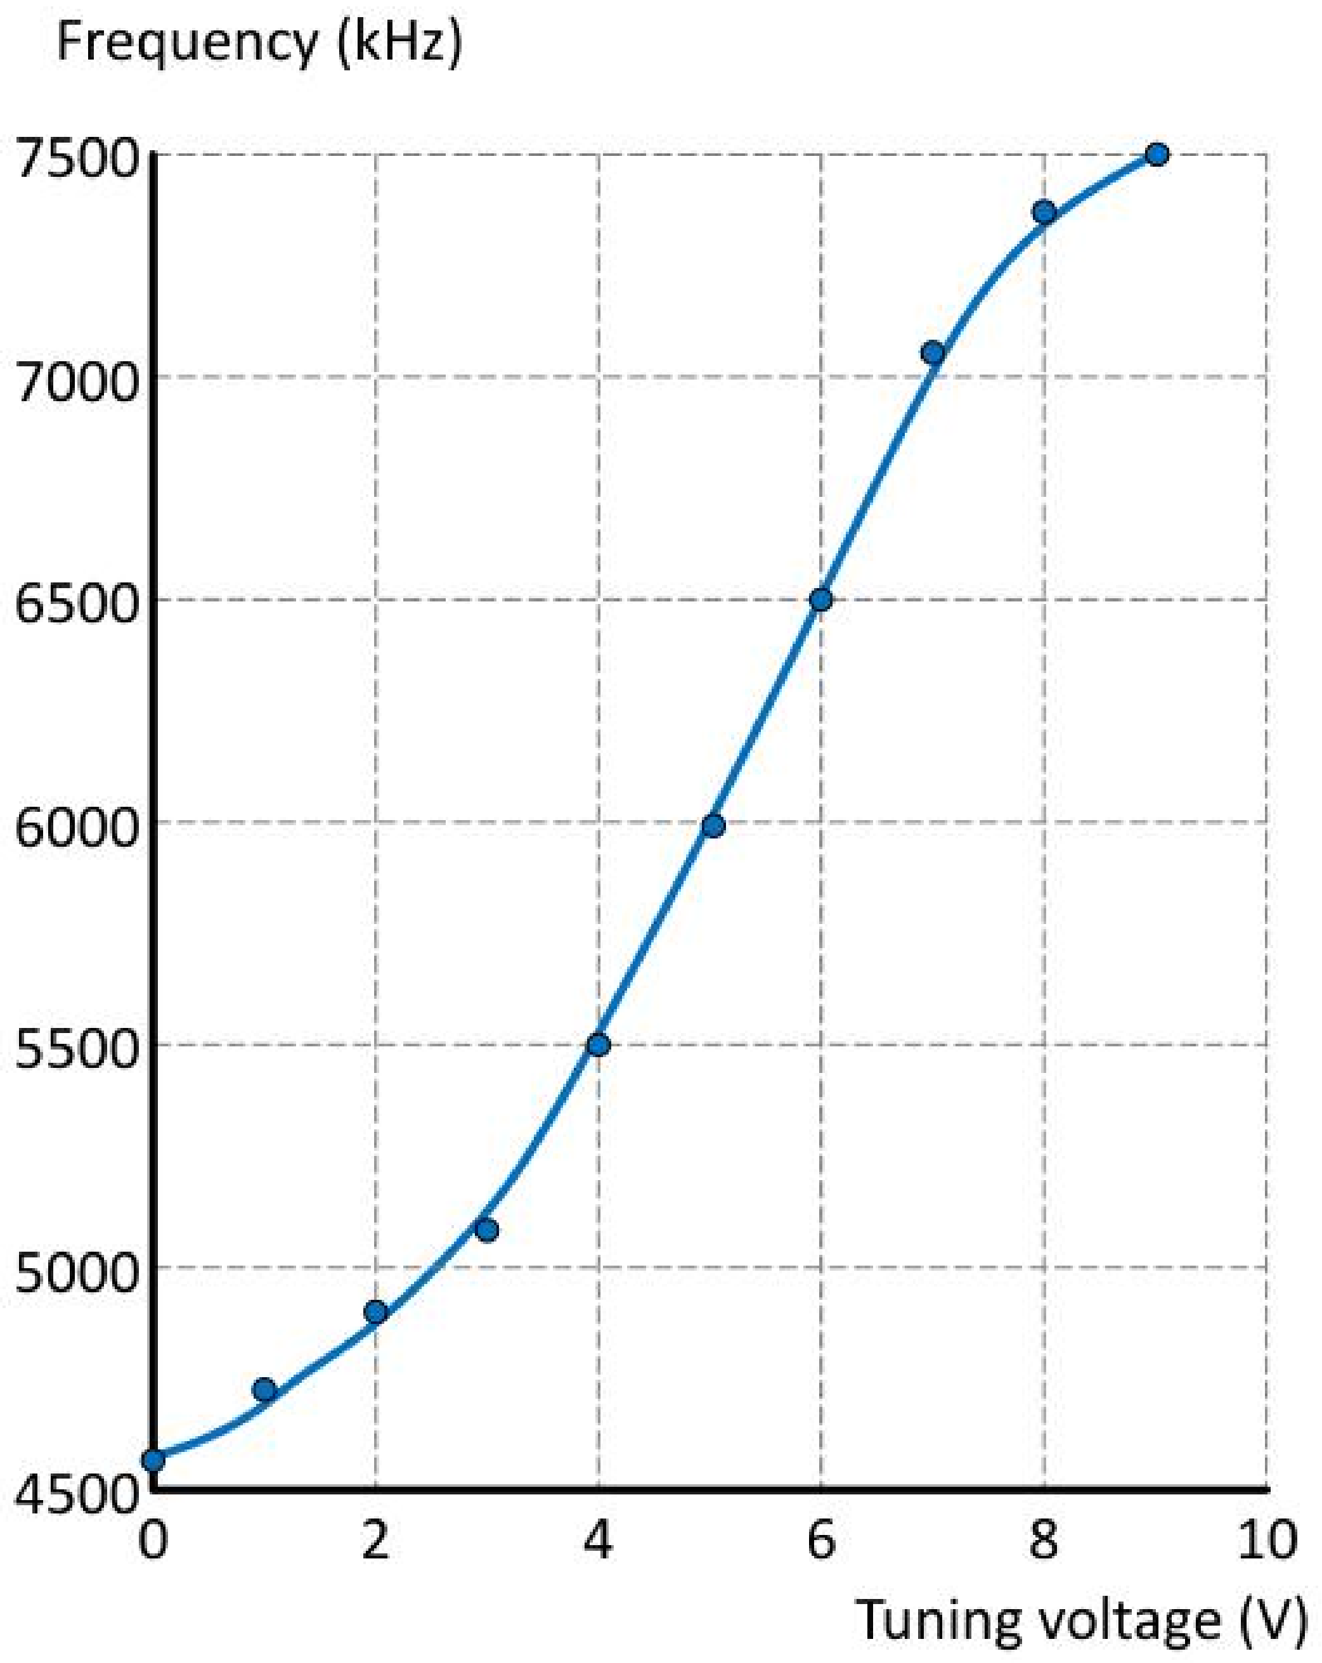

Fig.9.6: output frequency plotted against tuning voltage for the electronically tuned Colpitts oscillator shown in Fig.9.4.

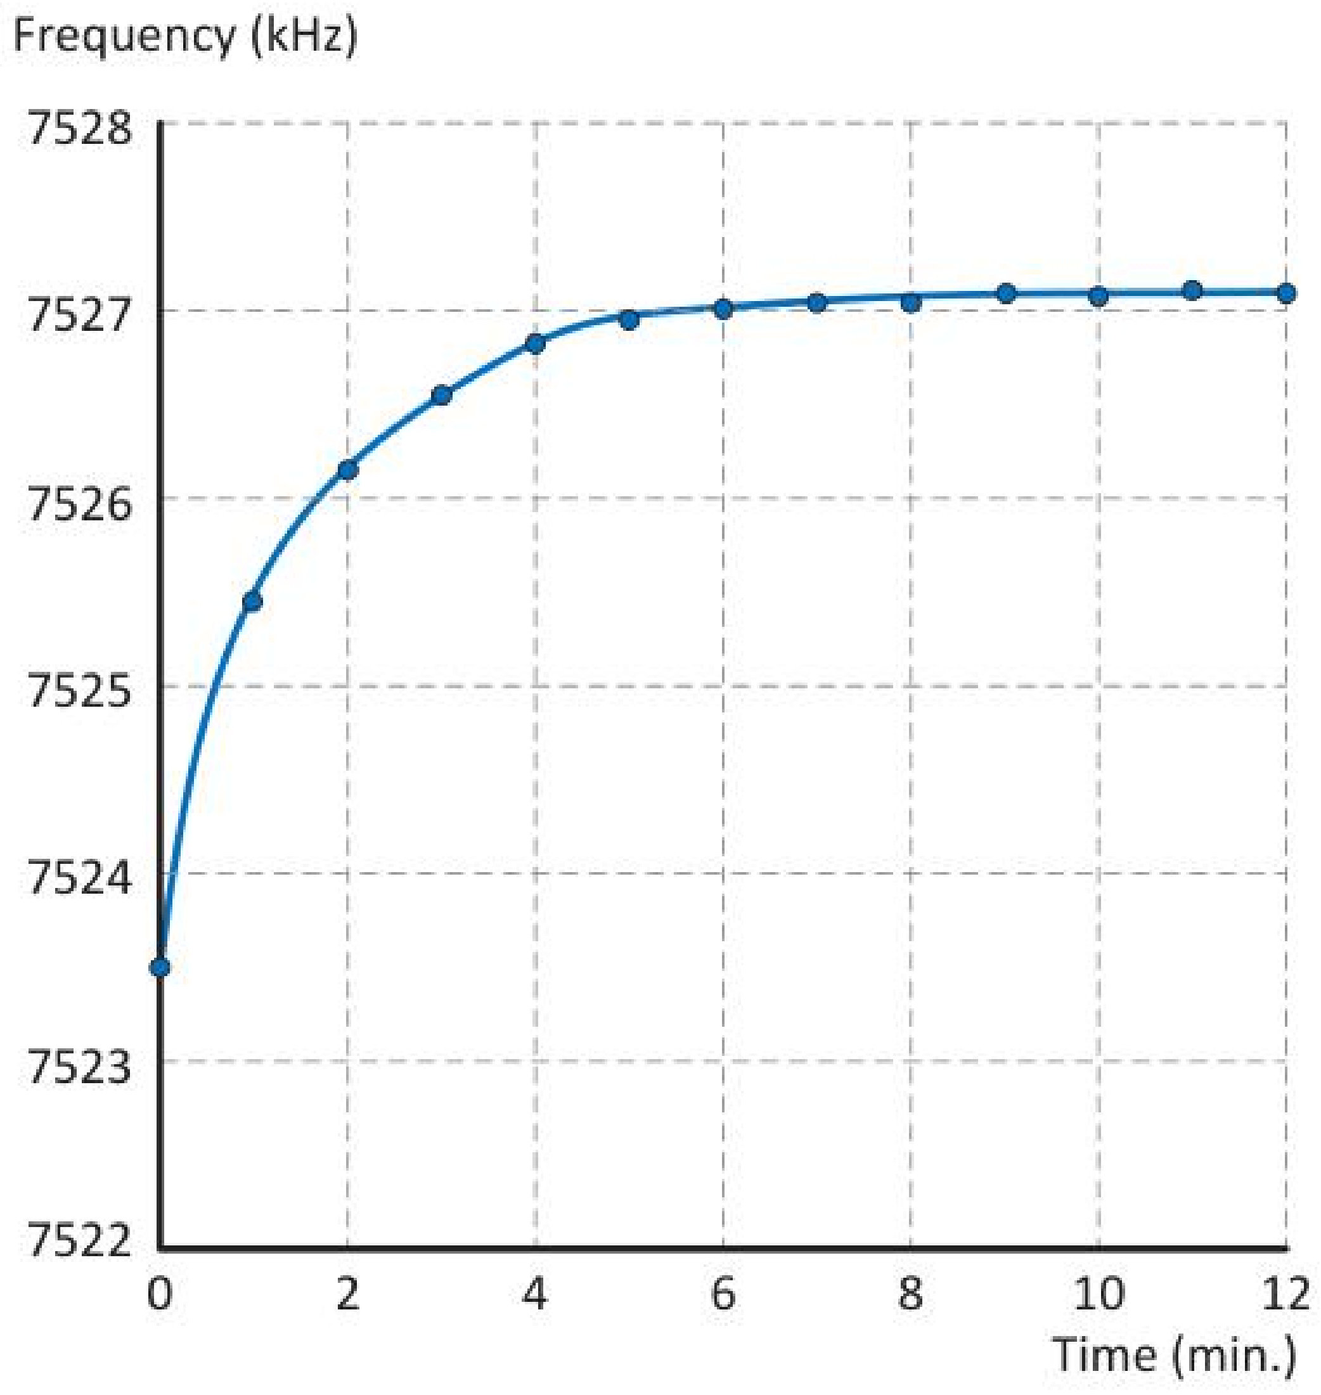

Fig.9.7: oscillator drift for the electronically tuned Colpitts oscillator shown in Fig.9.4.

D1 is operated in a reverse-biased state, with variable DC voltage from the tuning control, VR1. This takes advantage of the varicap's relatively large variation in capacitance with reverse-bias voltage. To promote frequency stability, the bias voltage is derived from a three-terminal 78L09 voltage regulator.

The 1SV149 variable-capacitance diode has a capacitance of about 450pF with an applied reverse bias of 1V DC, progressively falling to around 20pF when this reaches 9V. With the values in Fig.9.4, this variation in capacitance results in an overall 3MHz tuning range, from about 4.5MHz (4500kHz) to around 7.5MHz (7500kHz), as illustrated in Fig.9.6.

Linearity and frequency error

Fig.9.6 demonstrates the impact of non-linearity at both ends of the electronic tuning range of our modified Colpitts oscillator. However, if the tuning voltage can be limited to a narrower band, such as 4-7V, its linearity is significantly improved.

Linearity is the ratio between observed frequency error and overall tuning range, expressed as a percentage. Limiting the tuning range as suggested results in a frequency excursion of around 1.5MHz, with a significantly improved linearity of around ±3.3%.

Drift

Drift is another undesirable characteristic of all free-running oscillators. Drift is the variation of output frequency with time due to temperature changes, and is usually specified in Hz or kHz per unit time.

Drift is often characterised as either short-term or long-term. Short-term drift is invariably larger than long-term drift, resulting from the rapid increase in temperature during the brief period after power is initially applied to the oscillator circuit.

Short-term drift is therefore usually quoted for the first few minutes of operation. Thereafter, the output frequency tends to change at a much lower rate, and will usually stabilise at a value that remains reasonably constant.

Fig.9.7 shows the variation of output frequency for the Colpitts oscillator shown in Fig.9.4 when set initially to a frequency of 7523.5kHz. Note how the frequency increases by around 3.5kHz in the first five minutes, after which the rate of change of frequency progressively reduces, stabilising at around 7527kHz when the circuit reaches its working temperature.

Thereafter, the long-term drift is less than about 100Hz over a four-minute period (about 25Hz per minute). This is adequate for many non-critical applications, but the lack of frequency stabilisation makes it unsuitable for narrow-band operation and digital demodulation.

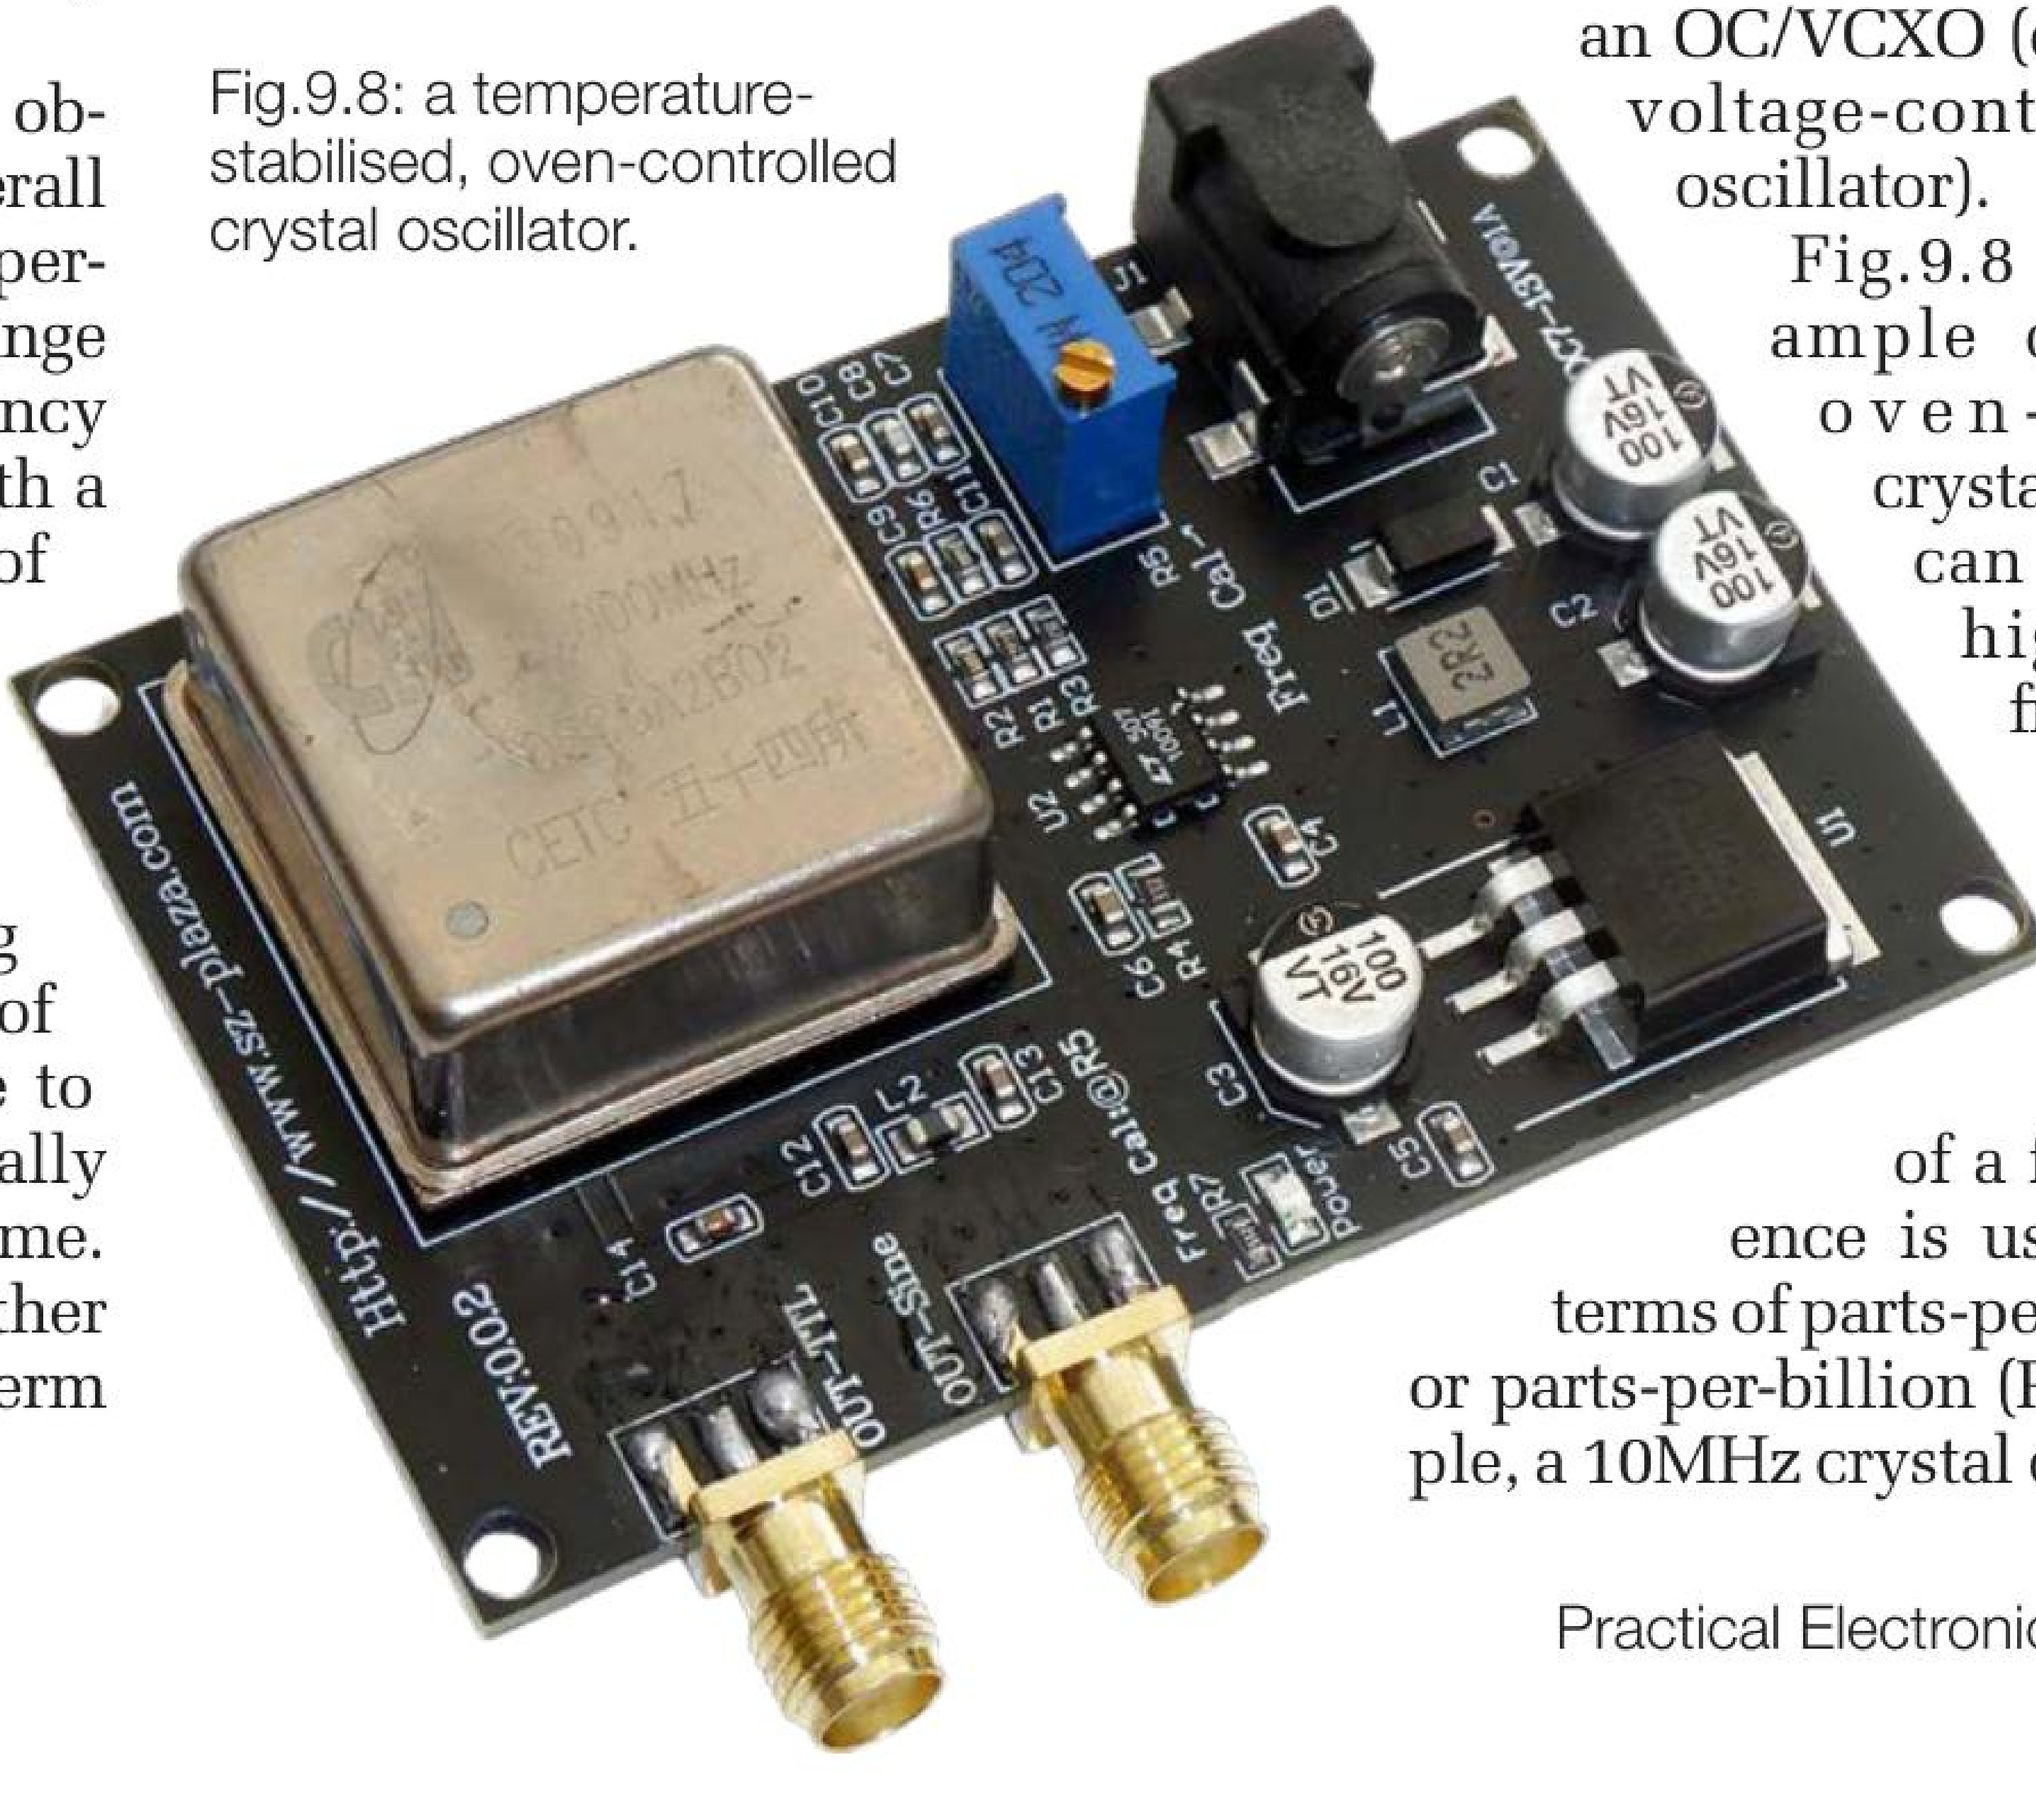

Fig.9.8: a temperature-stabilised, oven-controlled crystal oscillator.

Improved RF signal generating techniques

Because of their inherent non-linearity and a tendency to drift, free-running oscillators are avoided in critical RF applications. Instead, modern RF signal sources use phase-locked loops, frequency dividers and digital control methods, often controlled using an I^2C or SPI serial bus.

Various strategies are used for generating signals, some of which are based on a voltage-controlled crystal oscillator (VCXO) combined with a phase-locked loop and frequency divider. Others employ direct digital synthesis (DDS), with the signal generated digitally. We will be looking at both methods based on currently available low-cost modules.

An essential component in many frequency-generating systems is a variable crystal oscillator (VCXO). This comprises a crystal oscillator with limited tuning provided by a variable capacitance diode, enabling its output frequency to be adjusted electronically as described earlier.

To allow useful tuning around the nominal crystal frequency, a VCXO must trade some intrinsic stability for increased frequency pull range. This can degrade its frequency-versus-temperature characteristic, its ageing, its short-term stability and associated phase noise.

When a VCXO is used in a PLL, its frequency excursion must always be at least as great as the combined instability of the VCXO itself and the reference to which it is being locked.

A VCXO that incorporates additional circuitry to counter the effects of temperature change is referred to as a TC/VXO (temperature-compensated voltage-controlled oscillator). To further reduce the effects of temperature change, a small temperature-controlled oven can be added, creating an OC/VCXO (oven-controlled voltage-controlled crystal oscillator).

Fig.9.8 shows an example of a low-cost oven-controlled crystal oscillator that can be used as a highly accurate frequency reference.

Frequency stability

The stability of a frequency reference is usually stated in terms of parts-per-million (PPM) or parts-per-billion (PPB). For example, a 10MHz crystal oscillator with a

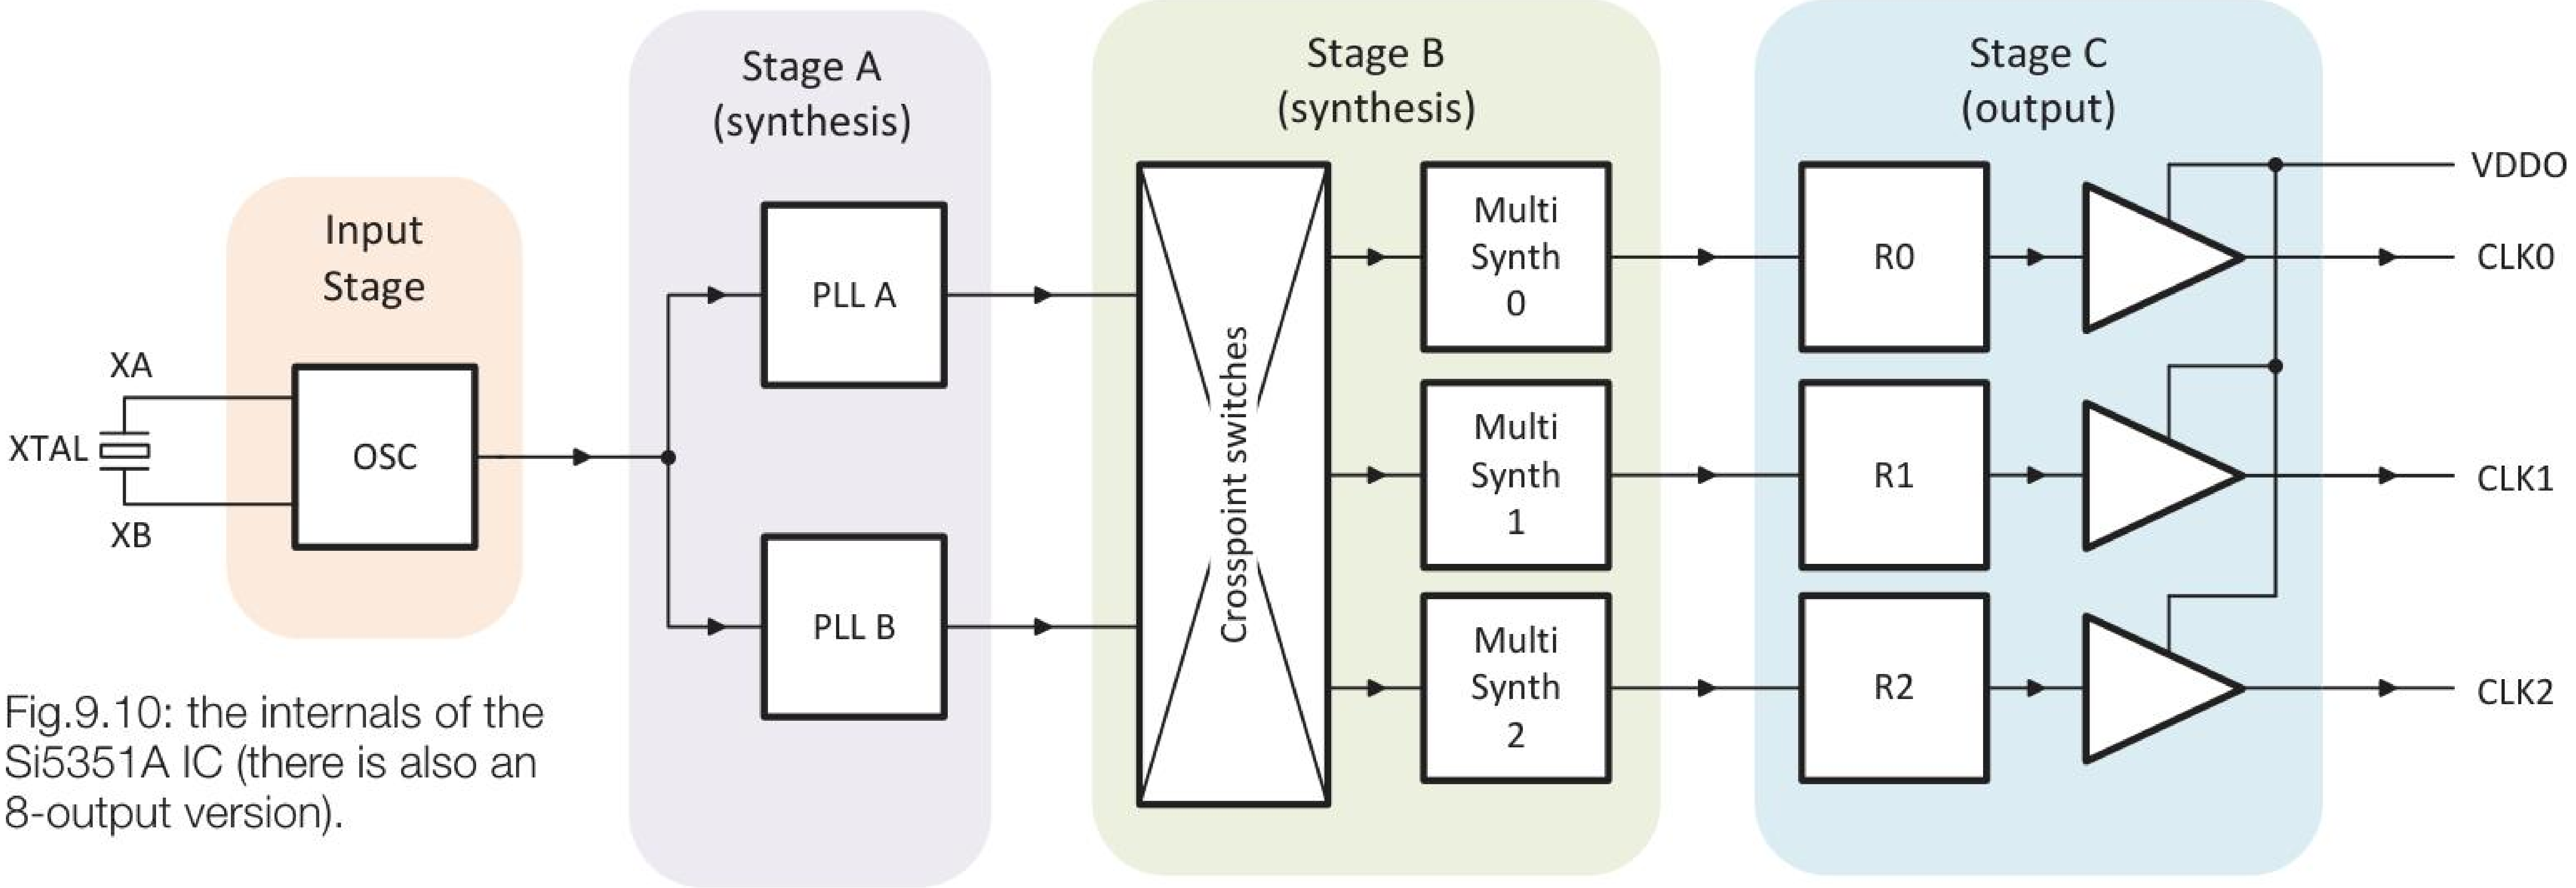

Fig.9.10: the internals of the Si5351A IC (there is also an 8-output version).

quoted stability of 10PPM will have a maximum frequency error of 100Hz.

The module shown in Fig.9.8 provides sine and square wave outputs at 10MHz with a quoted stability better than 20PPB. It consumes 600mA during its initial five-minute warm-up period but less than 250mA thereafter.

For accurate calibration, the output frequency can be adjusted by a control voltage obtained from a multi-turn preset potentiometer (R5). This derives its supply from a DC source that's closely regulated to ensure stability.

Next, we will move on to examine two contrasting approaches for accurate and stable frequency generation. The first is based on PLL techniques and frequency division; the second involves DDS techniques.

The Si5351 clock generator

The Si5351 is an I^{2}C-programmable frequency generator that's ideally suited for replacing one or more crystal-controlled oscillators. To aid the experimenter, the device is readily available in a variety of low-cost modules that can be easily incorporated into your own designs, as shown in Fig.9.9.

It's important to note that frequent references to the Si5351 as a DDS (direct digital synthesis) chip are incorrect because the device is based on PLL techniques and fractional dividers, rather than direct waveform synthesis used in DDS.

The Si5351 can generate precision frequencies up to 160MHz on each of its outputs. Three versions of the Si5351 are available:

- the Si5351A generates up to eight programmable clock frequencies from a single crystal-controlled master clock
- the Si5351B adds an internal VCXO and provides the flexibility to replace free-running clocks as well as synchronous clocks
- the Si5351C offers similar features but can also synchronise to an external reference clock (CLKIN)

Fig.9.10 shows a simplified schematic of the Si5351A. The device consists of an input stage, two synthesis stages (A and B) and an output stage (C). The input stage accepts an external crystal (XTAL) that sets the fundamental clock frequency.

The first synthesis stage (Stage A) multiplies the oscillator frequency using two phase-locked loops (PLL A and PLL B) to produce high-frequency intermediate clock signals. These two clocks (derived from Stage A) are then fed to a programmable cross-point switch within Stage B.

This flexible switching arrangement provides signal routing to the three 'MultiSynth' stages that act as high-resolution fractional dividers. The outputs from Stage B are fed to three further integer dividers, R0 to R2, at the input of Stage C. This additional stage of integer division allows the chip to generate output frequencies as low as 8kHz.

The divided frequency outputs are then fed to three buffer stages to produce the final output frequencies, CLK0 to CLK2.

Fig.9.11 (overleaf) shows a simplified schematic of the Si5351B. The device uses a similar arrangement to that of its smaller counterpart, but adds two further inputs: a clock input (CLKIN) and a control voltage input (VC). The former allows locking to an external clock input, while the latter provides frequency control from an external voltage control as well as the option of locking to an external clock input.

Stages B and C of the Si5351B are extended to provide eight outputs, CLK0 to CLK7, enabling the chip to generate free-running or synchronous non-integer related clock frequencies at each of its outputs. Note that it's possible for one device to synthesise clocks for multiple clock outputs.

The Si5351 uses a fixed-frequency standard AT-cut quartz crystal (introduced in part 2, December 2025) as a reference for its internal oscillator. The output of the oscillator can be used to provide a free-running reference for one or both of the phase-locked loops.

The oscillator requires a fundamental crystal frequency of either 25MHz or 27MHz; internal load capacitors are provided to avoid the need for external components. During software configuration, the capacitors can be set to provide a 6pF, 8pF or 10pF crystal load.

The Si5351B's external clock input is used as a clock reference for PLL A and PLL B; the external clock input (CLKIN) can accept any frequency between 10MHz and 100MHz.

The Si5351B's internal VCXO provides excellent linearity coupled with reliable startup and a wide tuning range (linear from 10% to 90% of VDD). A unique feature of the chip is its ability to generate multiple output frequencies determined by the same control voltage. The chip can thus replace multiple PLLs or VCXOs that would normally be locked to a common reference.

By applying a control voltage (Vcontrol), the device can function without an external pullable quartz crystal.

As with the Si5351B, the Si5351A uses two synthesis stages (A and B) to generate its final RF output from

Fig.9.9: a low-cost clock generator module based on the Si5351A chip.

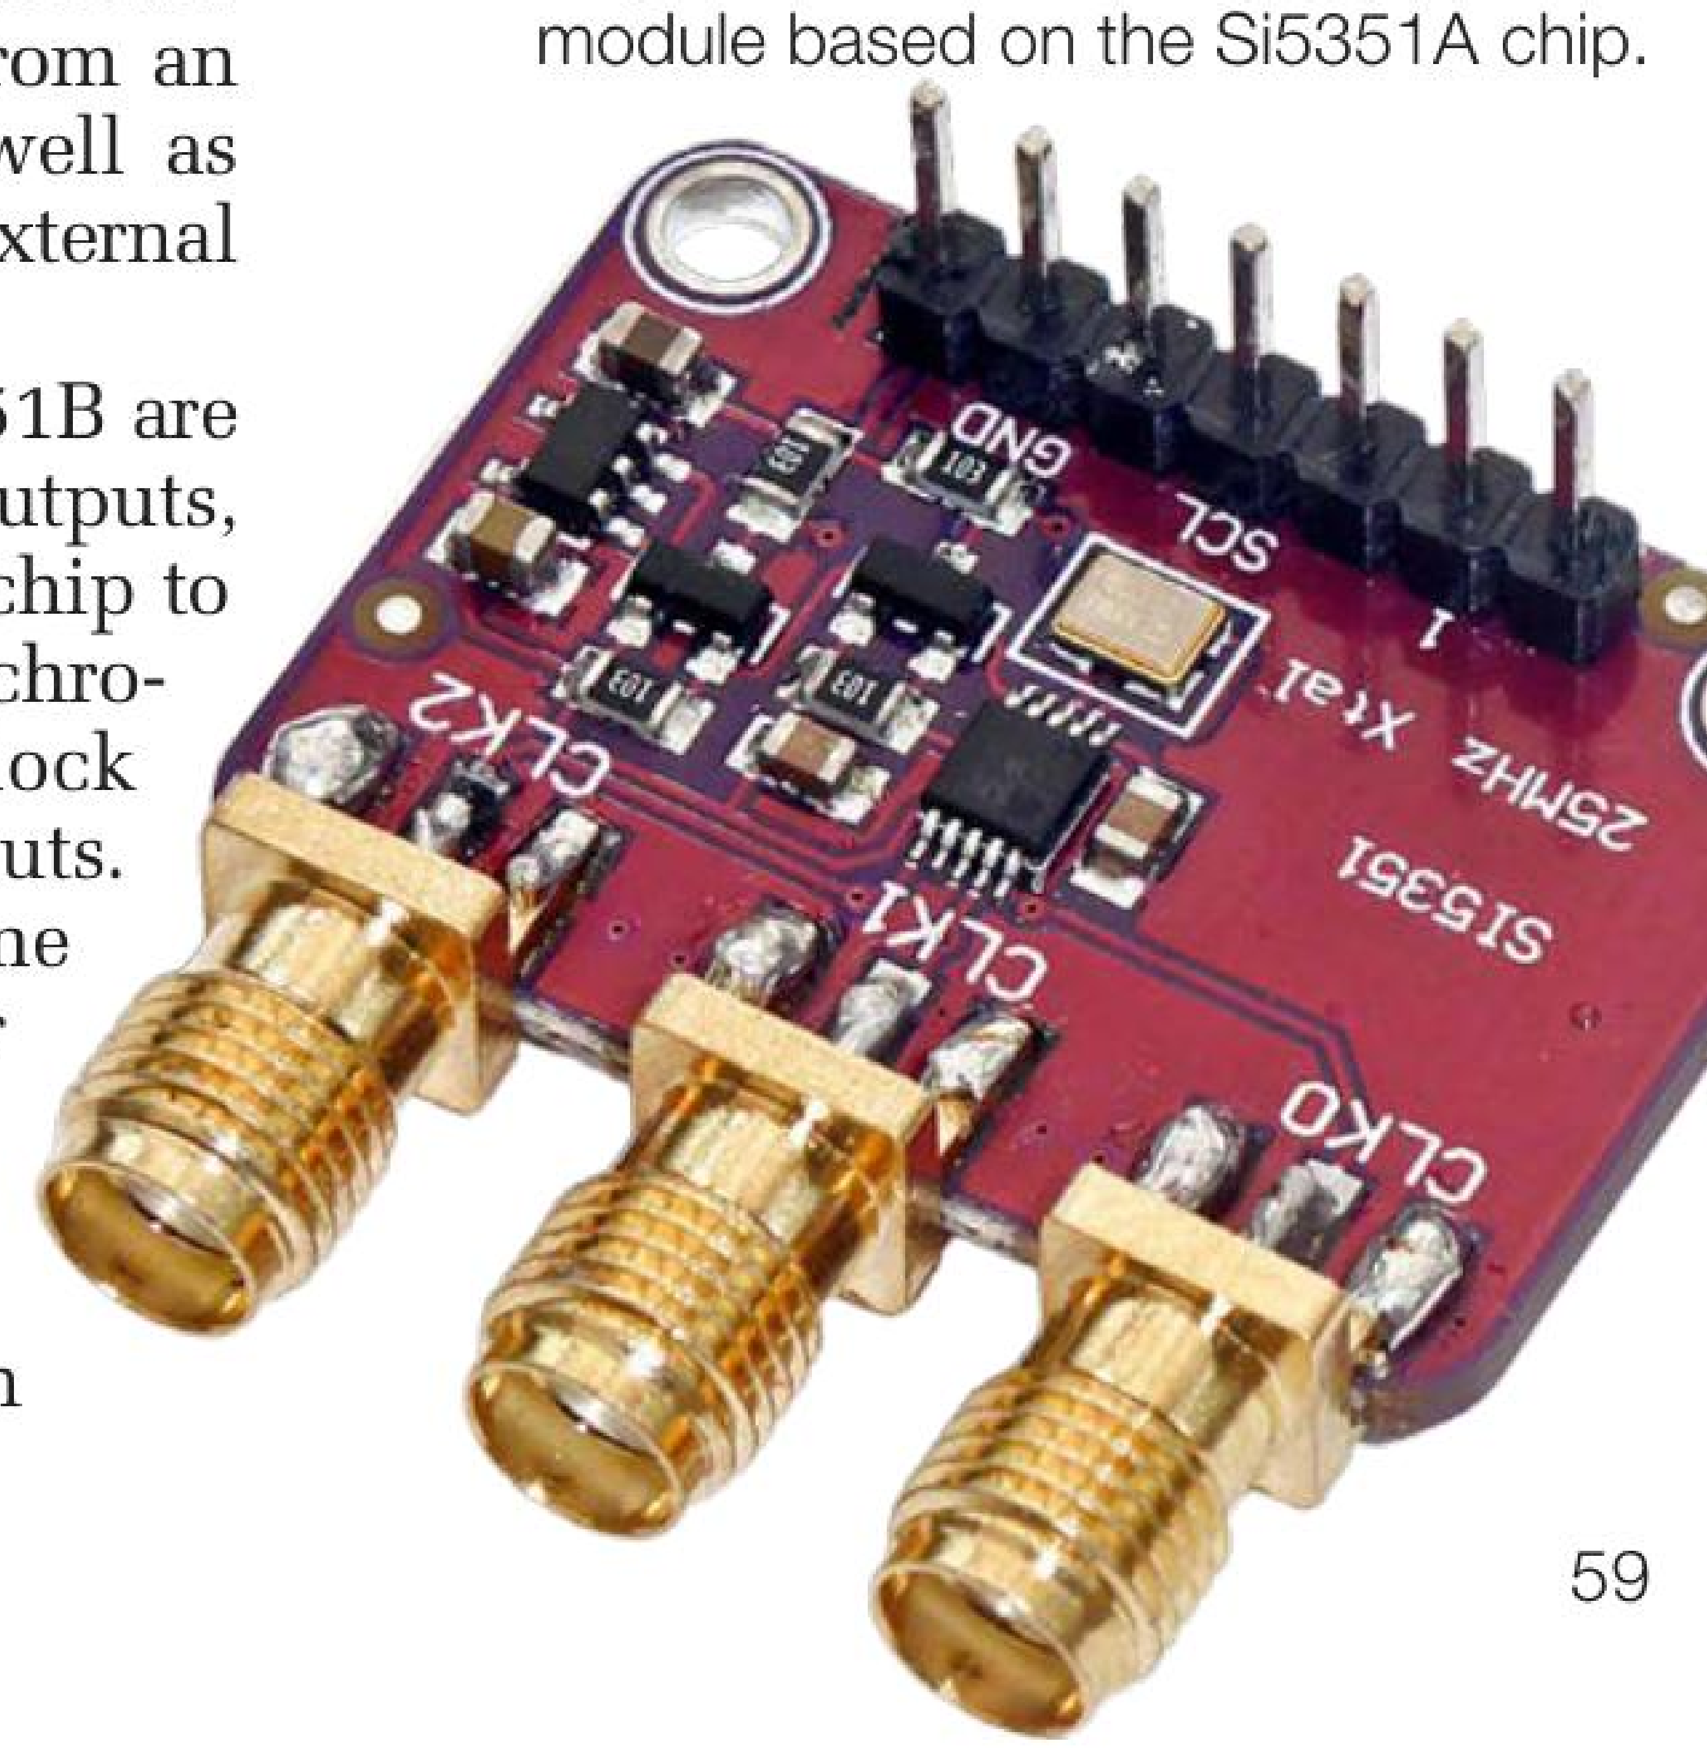

Fig.9.11: the internal arrangement of the Si5351B IC. This version of the chip always has eight outputs.

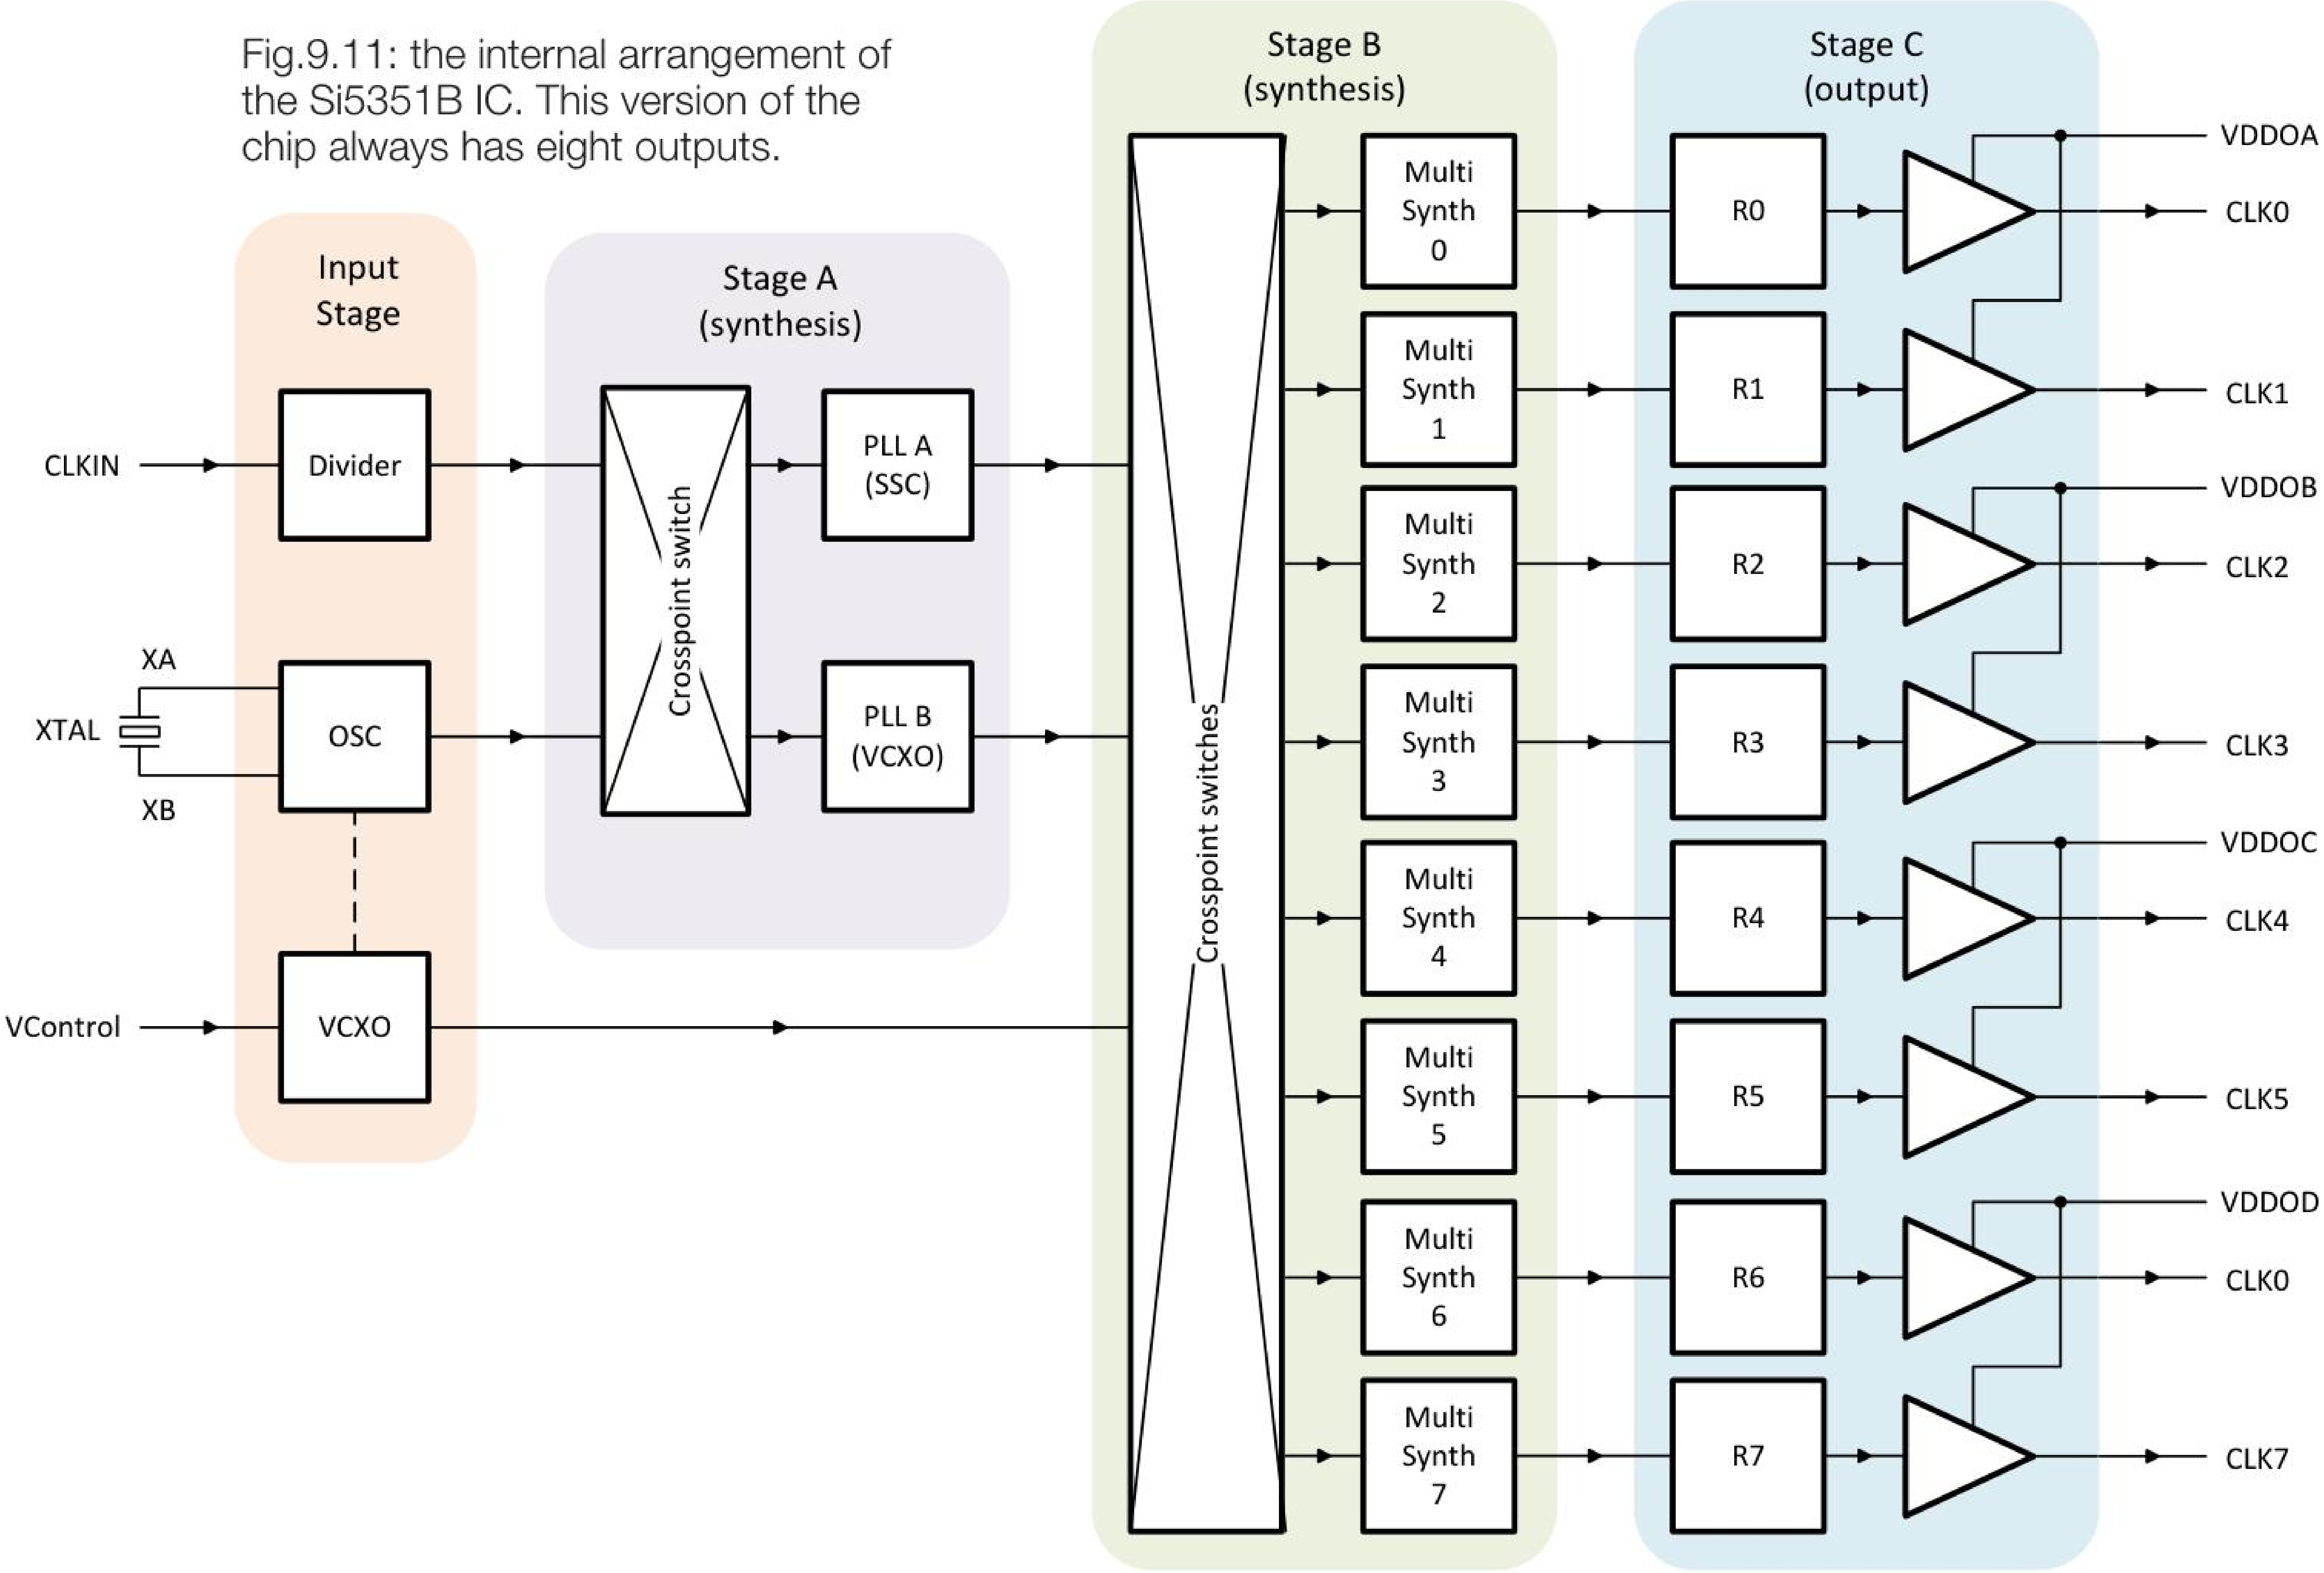
Input Stage
Stage A (synthesis)
Stage B (synthesis)
Stage C (output)
CLKIN
Divider
XA
XTAL
XB
OSC
VControl
VCXO
Crosspoint switch
PLL A (SSC)
PLL B (VCXO)
Crosspoint switches
Multi Synth 0
Multi Synth 1
Multi Synth 2
Multi Synth 3
Multi Synth 4
Multi Synth 5
Multi Synth 6
Multi Synth 7
R0
R1
R2
R3
R4
R5
R6
R7
VDDOA
CLK0
CLK1
VDDOB
CLK2
CLK3
VDDOC
CLK4
CLK5
VDDOD
CLK0
CLK7

CLK0 to CLK7. Stage A uses PLLs to multiply the lower frequency input references to a high-frequency intermediate clock.

The second stage uses high-resolution fractional dividers provided by MultiSynth 0 to MultiSynth 7. Frequencies of up to 160MHz can be generated on multiple outputs.

The Stage A (input) cross-point switch enables each PLL to lock onto either the CLKIN or XTAL input. This allows the PLLs to use the same or different sources for independent or synchronous clock generation.

Stage B's cross-point switch lets any fractional divider receive input from either PLL A or PLL B, so outputs can be synchronous or non-synchronous at unrelated, non-integer frequencies as required.

The VCXO in the Si5351B already generates a high-frequency intermediate clock, so it can be fed directly into Stage B. The high-resolution fractional dividers allow synthesis of the VCXO centre frequency from 400kHz to 160MHz.

Frequency generation with the Si5351

It's worth providing a few simple examples of frequency generation with this chip. Fig.9.12 illustrates the various stages of conversion, division and synthesis within the device.

The output of Stage A (f_{VCO}) is a multiple of the crystal oscillator frequency (f_{XO}). The multiple is determined by the Stage A fractional division ratio, Div.A. In turn, Div.A is determined by variables a, b and c. The yellow note in Fig.9.12 shows the range of values that can be used for these.

The output of Stage B (f) is a ratio of the VCO frequency (f_{VCO}). The ratio is determined by the Stage B fractional division ratio, Div.B. In turn, Div.B is determined by variables d, e and f. As before, the yellow note in Fig.9.12 shows the range of values that can be used for these quantities.

The output of Stage C (f_{out}) is a ratio

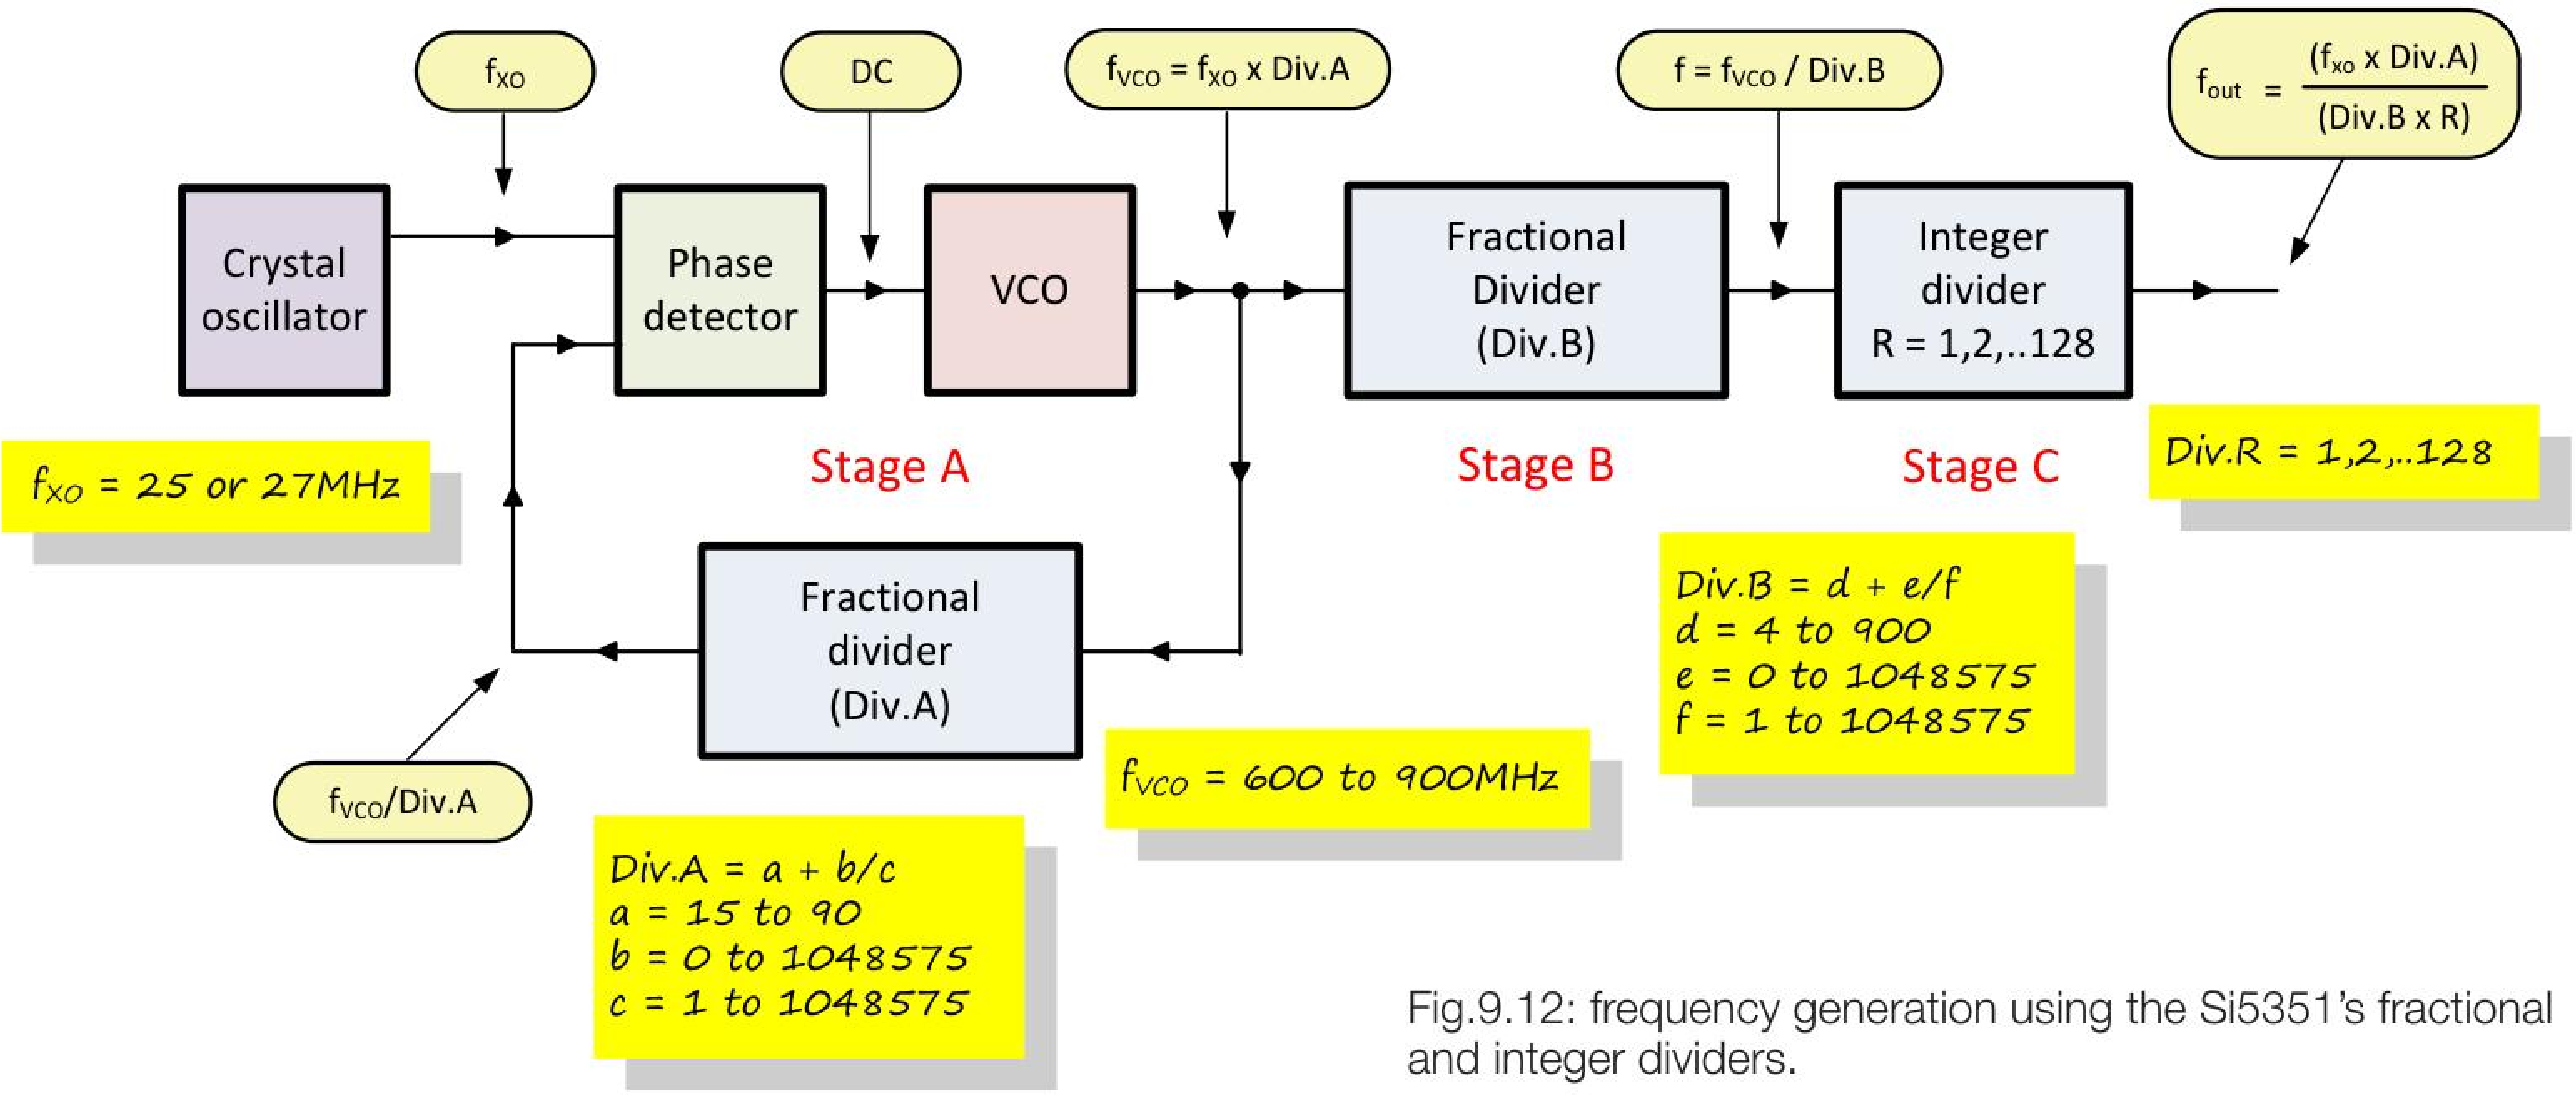
f_{XO}
DC
$f_{VCO} = f_{XO}$ x Div.A
$f = f_{VCO}$ / Div.B
$f_{out} = \frac{(f_{xo} \text{ x Div.A})}{(\text{Div.B x R})}$
Crystal oscillator
Phase detector
VCO
Fractional Divider (Div.B)
Integer divider R = 1,2,..128
f_{XO} = 25 or 27MHz
Stage A
Stage B
Stage C
Div.R = 1,2,..128
Fractional divider (Div.A)
Div.B = d + e/f
d = 4 to 900
e = 0 to 1048575
f = 1 to 1048575
f_{VCO}/Div.A
f_{VCO} = 600 to 900MHz
Div.A = a + b/c
a = 15 to 90
b = 0 to 1048575
c = 1 to 1048575

Fig.9.12: frequency generation using the Si5351's fractional and integer dividers.

of the output frequency from Stage B (*f*). The ratio is determined by the Stage C division ratio, *R*, which can take an integer value from 1 to 128. We can now combine these relationships to obtain the following.

The output frequency, f_{VCO}, from Stage A is given by:

$$f_{vco} = f_{xo} \cdot \text{Div.A} = f_{xo} \cdot \left(a + \frac{b}{c}\right)$$

Here, *a*, *b* and *c* determine the division ratio provided by the Stage A fractional divider.

The output frequency, *f*, from Stage B is given by:

$$f = \frac{f_{vco}}{\text{Div.B}} = \frac{f_{vco}}{\left(d + \frac{e}{f}\right)}$$

Here, *d*, *e* and *f* determine the division ratio provided by the Stage B fractional divider.

The third (and final) stage of synthesis applies an additional integer ratio of division, *R*, in Stage C. The final output frequency can then be determined from:

$$f_{out} = \frac{f_{xo} \cdot \text{Div.A}}{\text{Div.B} \cdot R} = \frac{f_{xo} \cdot \left(a + \frac{b}{c}\right)}{\left(d + \frac{e}{f}\right) \cdot R}$$

Fig.9.13 shows how the Si5351 can generate outputs at (a) 55MHz and (b) 55kHz. In the second case, the a fractional division ratio is used for Div.A and the additional integer divider applied by Div.R.

To simplify the design process when using the Si5351 series of chips, you will find the Skyworks ClockBuilder Pro. utility extremely useful. This software will guide you through each of the stages of a clock design, as shown in Figs.9.14 & 9.15.

Programming the Si5351

Programming the Si5351 is simple and can be done using any low-cost microcontroller that supports the I^{2}C bus. An example Arduino code fragment is shown in Listing 9.1 (overleaf). This can be used to generate a constant 10MHz square wave output from the Si5351's CLK0 output at a level of approximately 10dBm (10mW).

An Si5351 module, like that shown in Fig.9.8, can be driven from the Arduino's I^{2}C data and clock signals (pins A4 and A5 on an Arduino Nano). An additional ground connection is also required.

The code is greatly simplified by the use of two library files, "si5351.h" and "Wire.h" (the latter is used for I^{2}C communication). If you need a

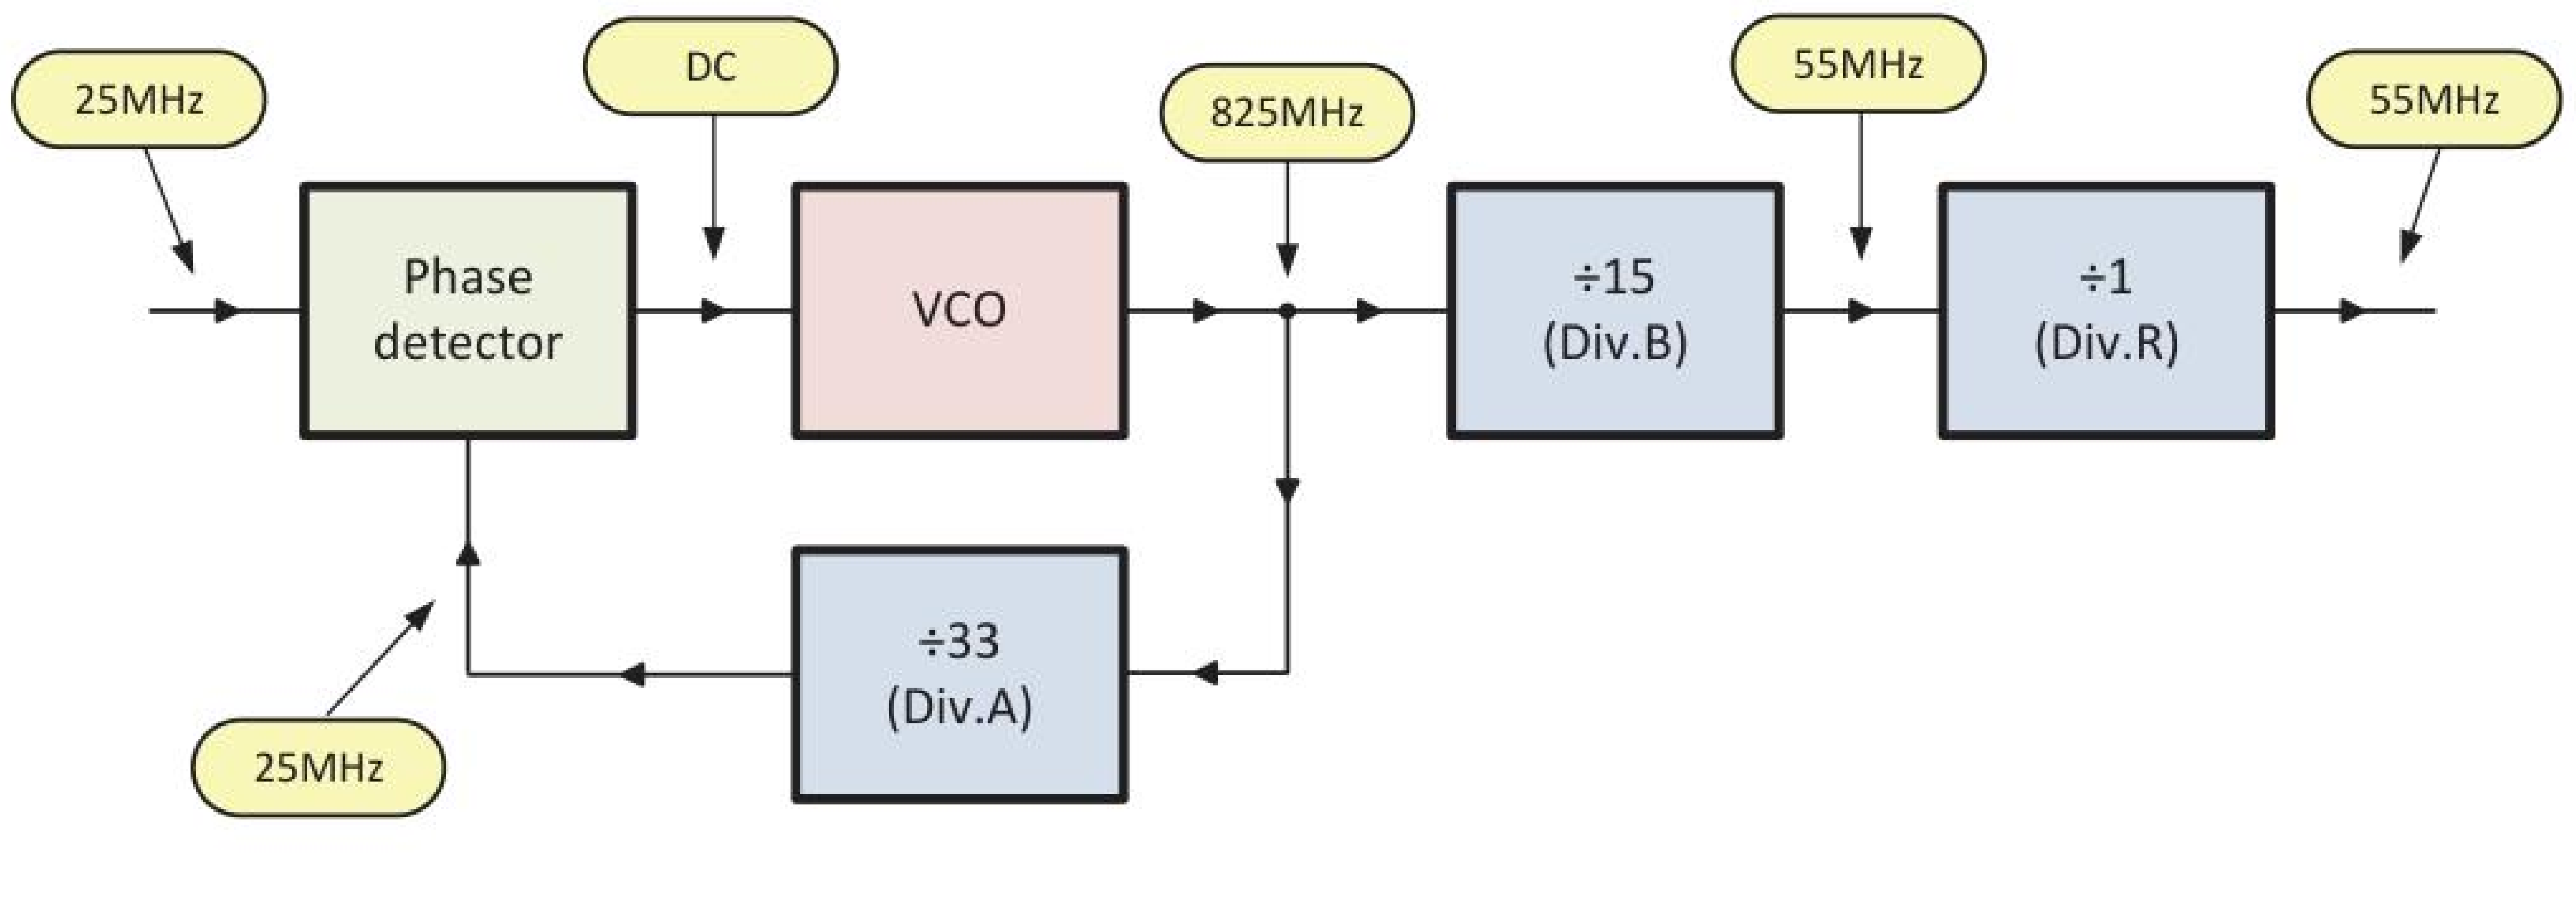
25MHz
DC
825MHz
55MHz
55MHz
Phase detector
VCO
÷15 (Div.B)
÷1 (Div.R)
÷33 (Div.A)
25MHz

(a) 55MHz

(b) 55kHz

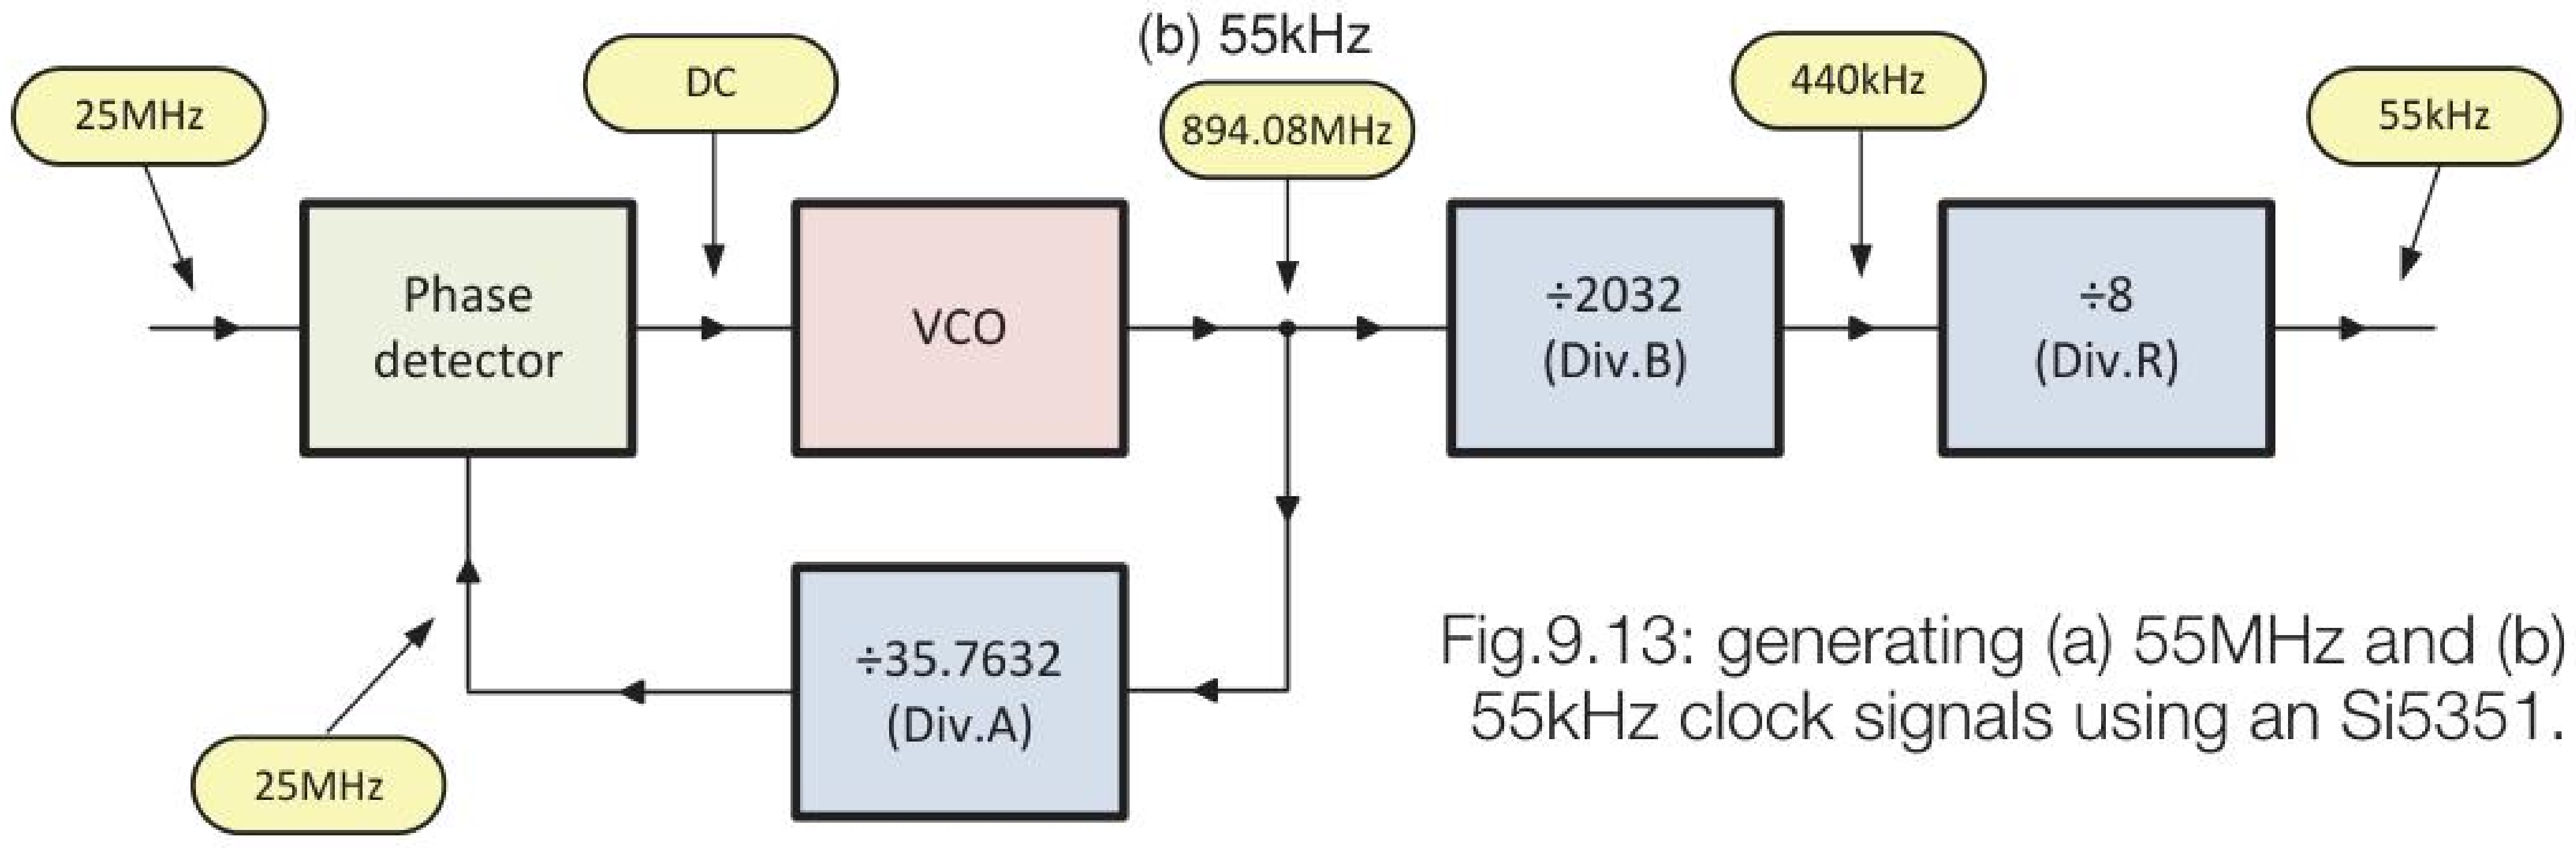
25MHz
DC
894.08MHz
440kHz
55kHz
Phase detector
VCO
÷2032 (Div.B)
÷8 (Div.R)
÷35.7632 (Div.A)
25MHz

Fig.9.13: generating (a) 55MHz and (b) 55kHz clock signals using an Si5351.

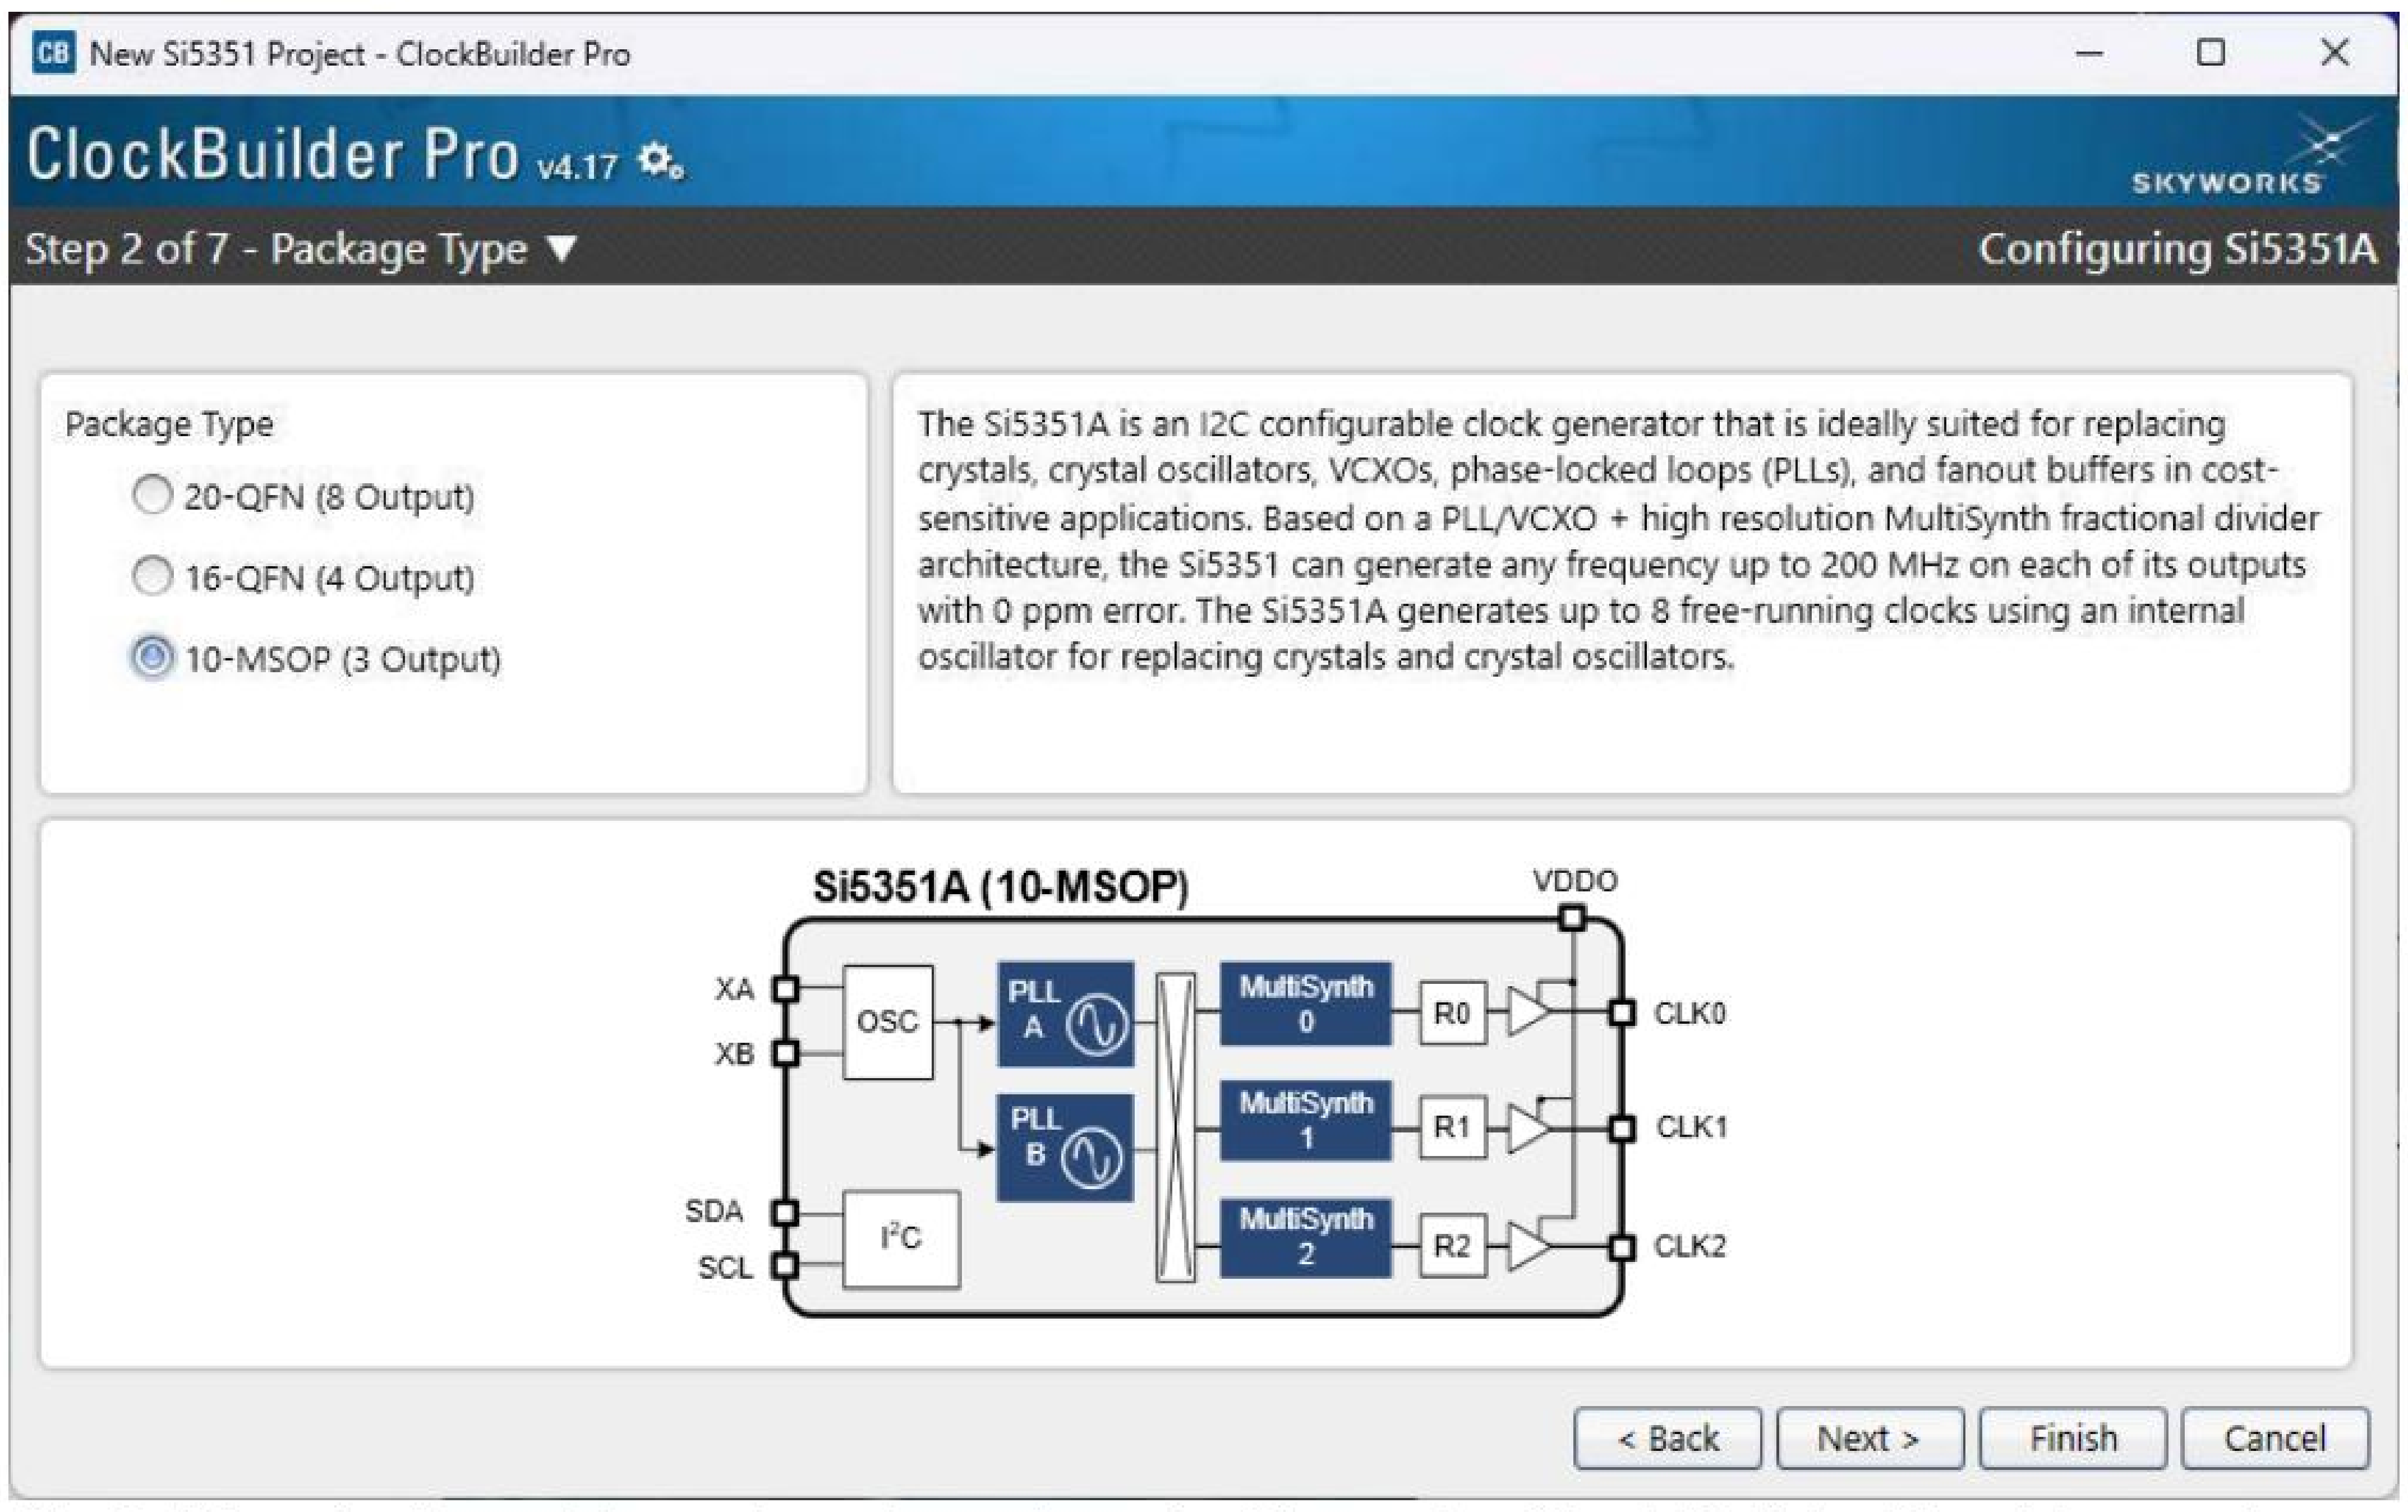

New Si5351 Project - ClockBuilder Pro

ClockBuilder Pro v4.17

Step 2 of 7 - Package Type ▼

Configuring Si5351A

Package Type

○ 20-QFN (8 Output)

○ 16-QFN (4 Output)

● 10-MSOP (3 Output)

The Si5351A is an I2C configurable clock generator that is ideally suited for replacing crystals, crystal oscillators, VCXOs, phase-locked loops (PLLs), and fanout buffers in cost-sensitive applications. Based on a PLL/VCXO + high resolution MultiSynth fractional divider architecture, the Si5351 can generate any frequency up to 200 MHz on each of its outputs with 0 ppm error. The Si5351A generates up to 8 free-running clocks using an internal oscillator for replacing crystals and crystal oscillators.

SKYWORKS
Si5351A (10-MSOP)
VDDO
XA
XB
OSC
PLL A
PLL B
SDA
SCL
I²C
MultiSynth 0
MultiSynth 1
MultiSynth 2
R0
R1
R2
CLK0
CLK1
CLK2
< Back
Next >
Finish
Cancel

Fig.9.14: selecting chip and package type in Skyworks ClockBuilder Pro. Here we've selected the 10-pin MSOP packaged version of the Si5351A.

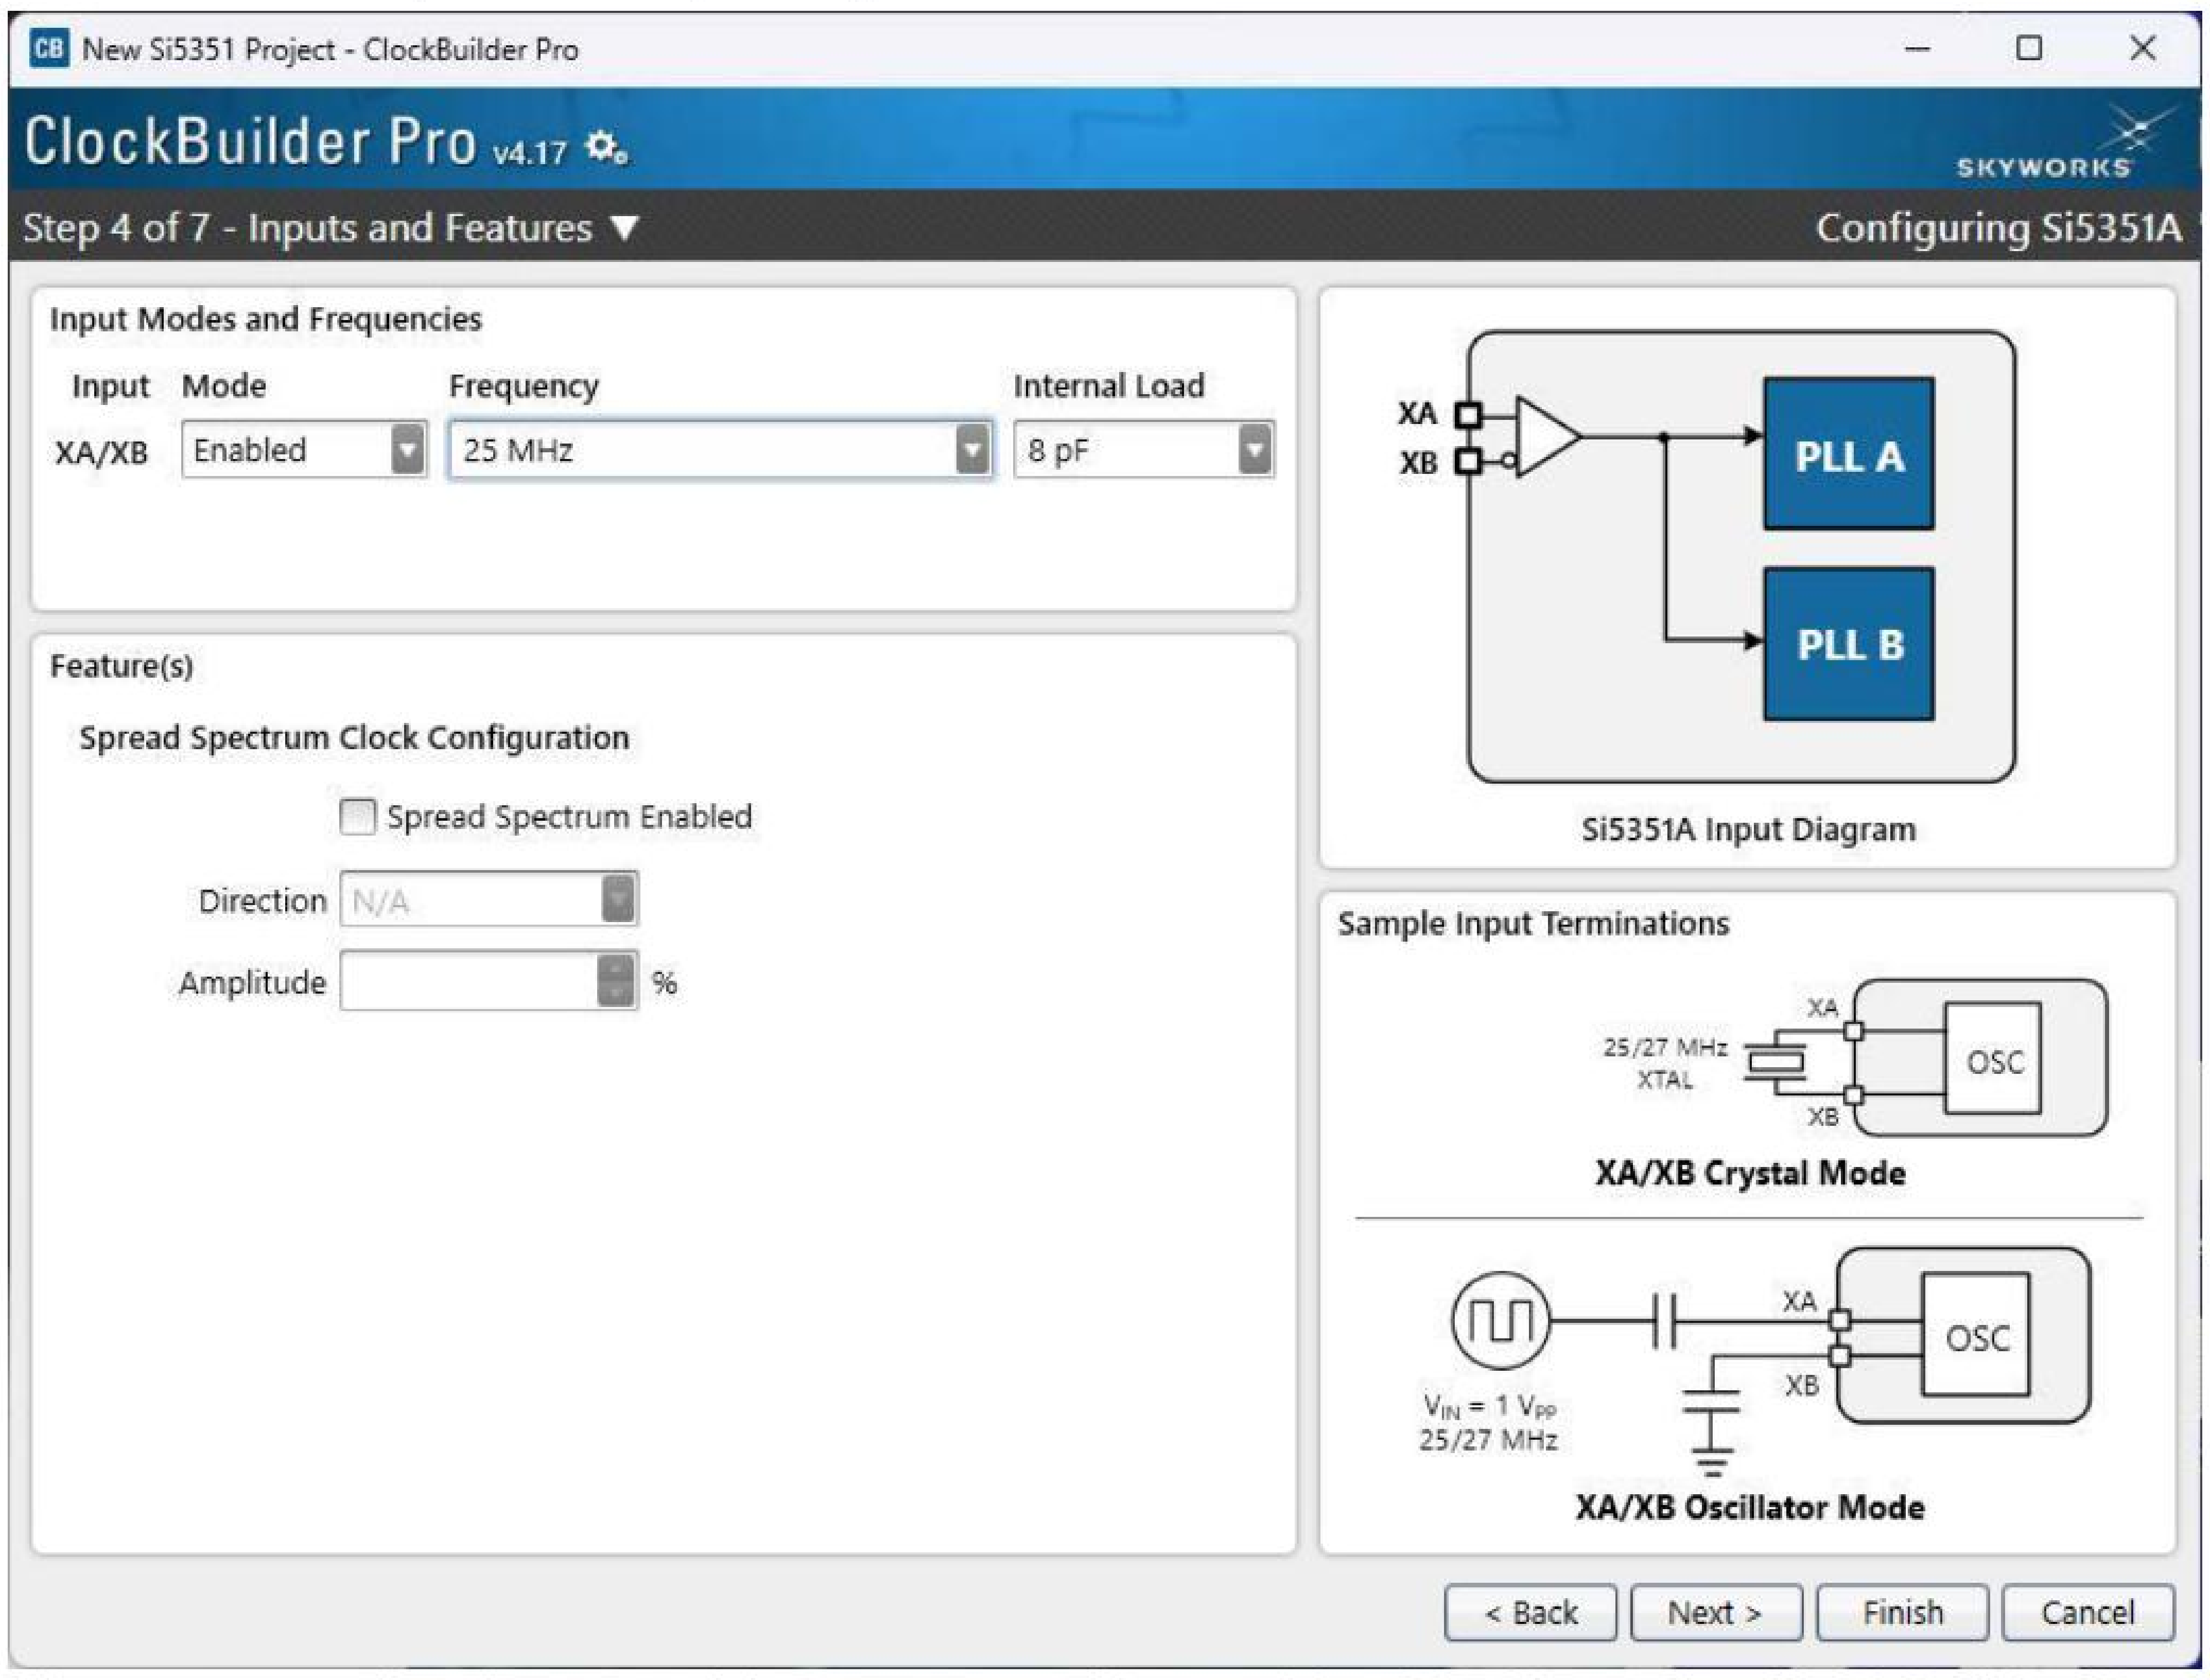

New Si5351 Project - ClockBuilder Pro

ClockBuilder Pro v4.17

Step 4 of 7 - Inputs and Features ▼

Configuring Si5351A

Input Modes and Frequencies

Input	Mode	Frequency	Internal Load
XA/XB	Enabled	25 MHz	8 pF

Feature(s)

Spread Spectrum Clock Configuration

☐ Spread Spectrum Enabled

Direction N/A

Amplitude %

SKYWORKS
XA
XB
PLL A
PLL B
Si5351A Input Diagram
Sample Input Terminations
25/27 MHz XTAL
XA
XB
OSC
XA/XB Crystal Mode
V_IN = 1 V_PP 25/27 MHz
XA
XB
OSC
XA/XB Oscillator Mode
< Back
Next >
Finish
Cancel

Fig.9.15: entering the crystal frequency and internal load in Skyworks ClockBuilder Pro.

Listing 9.1: an example code fragment for the Si5351 to produce a 10MHz, 10dBm output.

```
#include "si5351.h"
#include "Wire.h"

Si5351 si5351;

void setup() {
  bool i2c_found;
  // Start serial and initialize the Si5351
  Serial.begin(57600);
  i2c_found = si5351.init(SI5351_CRYSTAL_LOAD_8PF, 0, 0);
  if(!i2c_found)
    Serial.println("Device not found!");

  // Set CLK0 to 10 MHz and 10dBm appx.
  si5351.set_freq(1000000000ULL, SI5351_CLK0);
  si5351.drive_strength(SI5351_CLK0, SI5351_DRIVE_8MA);

  Serial.println("CLK0 set to 10MHz at 10dBm");

  // Query status update
  si5351.update_status();
  delay(500);
}

void loop() {
  // Other code goes here
}
```

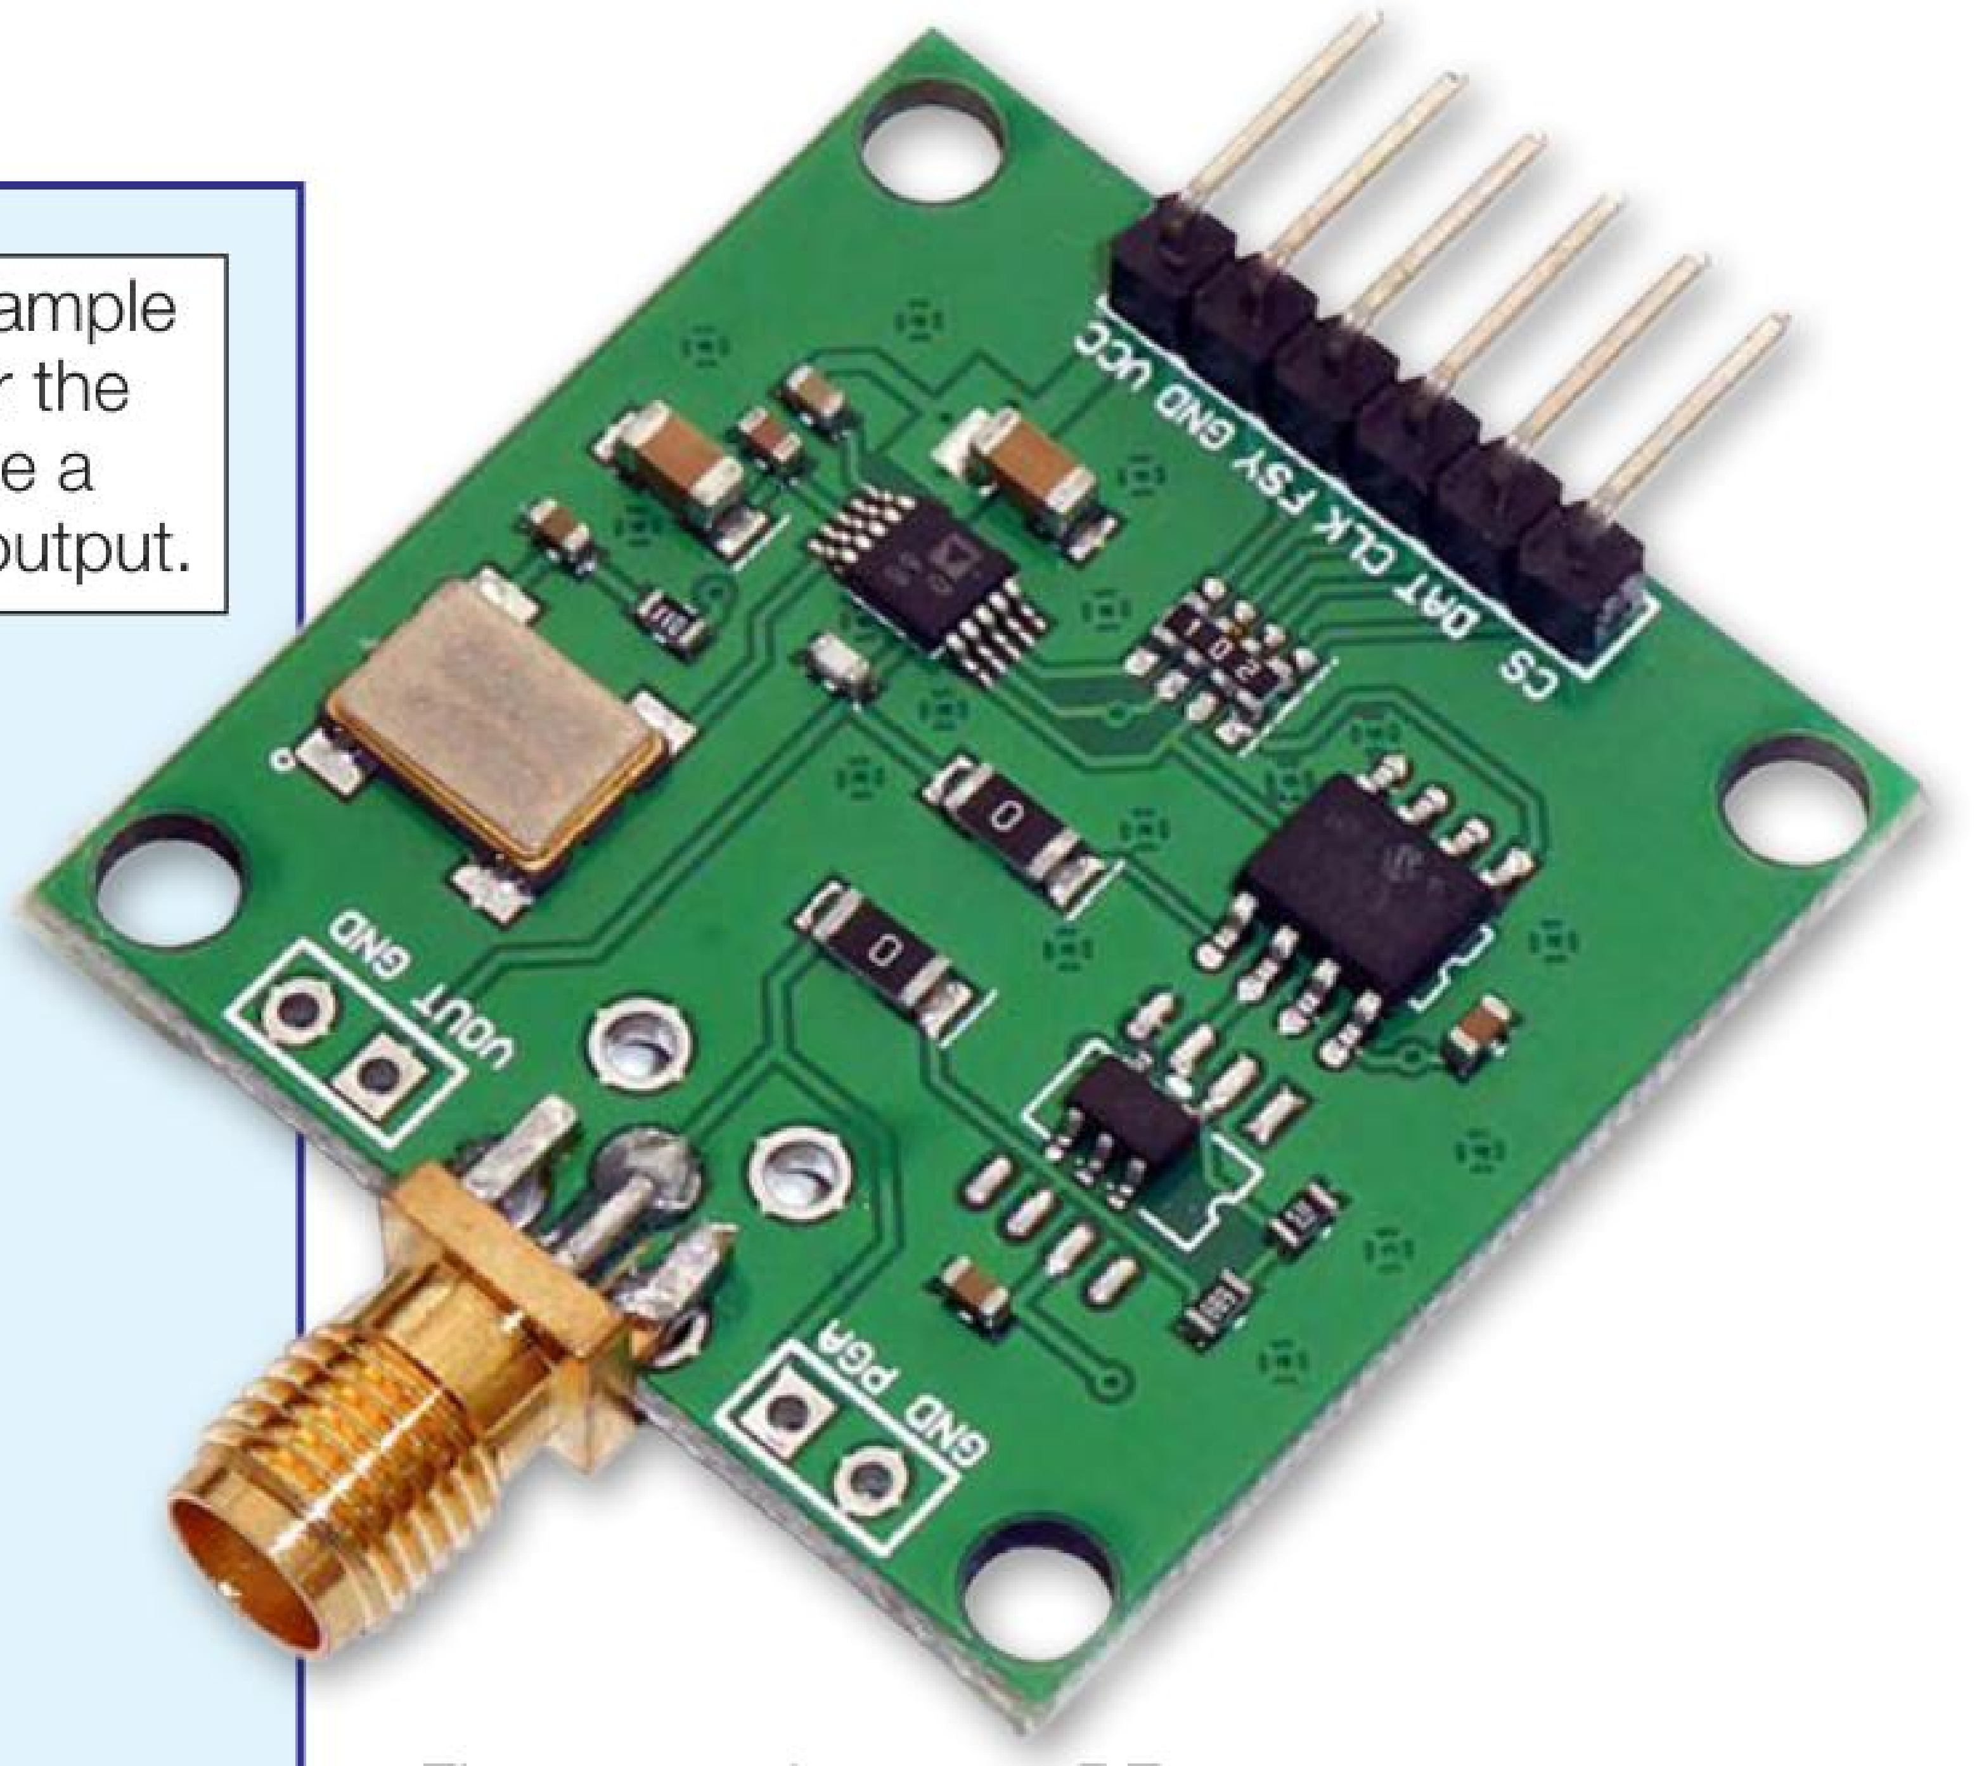
CS DAT CLK FSY GND VCC
VOUT GND
PGA GND

Fig.9.16: a low-cost RF generator module based on the AD9833 DDS chip.

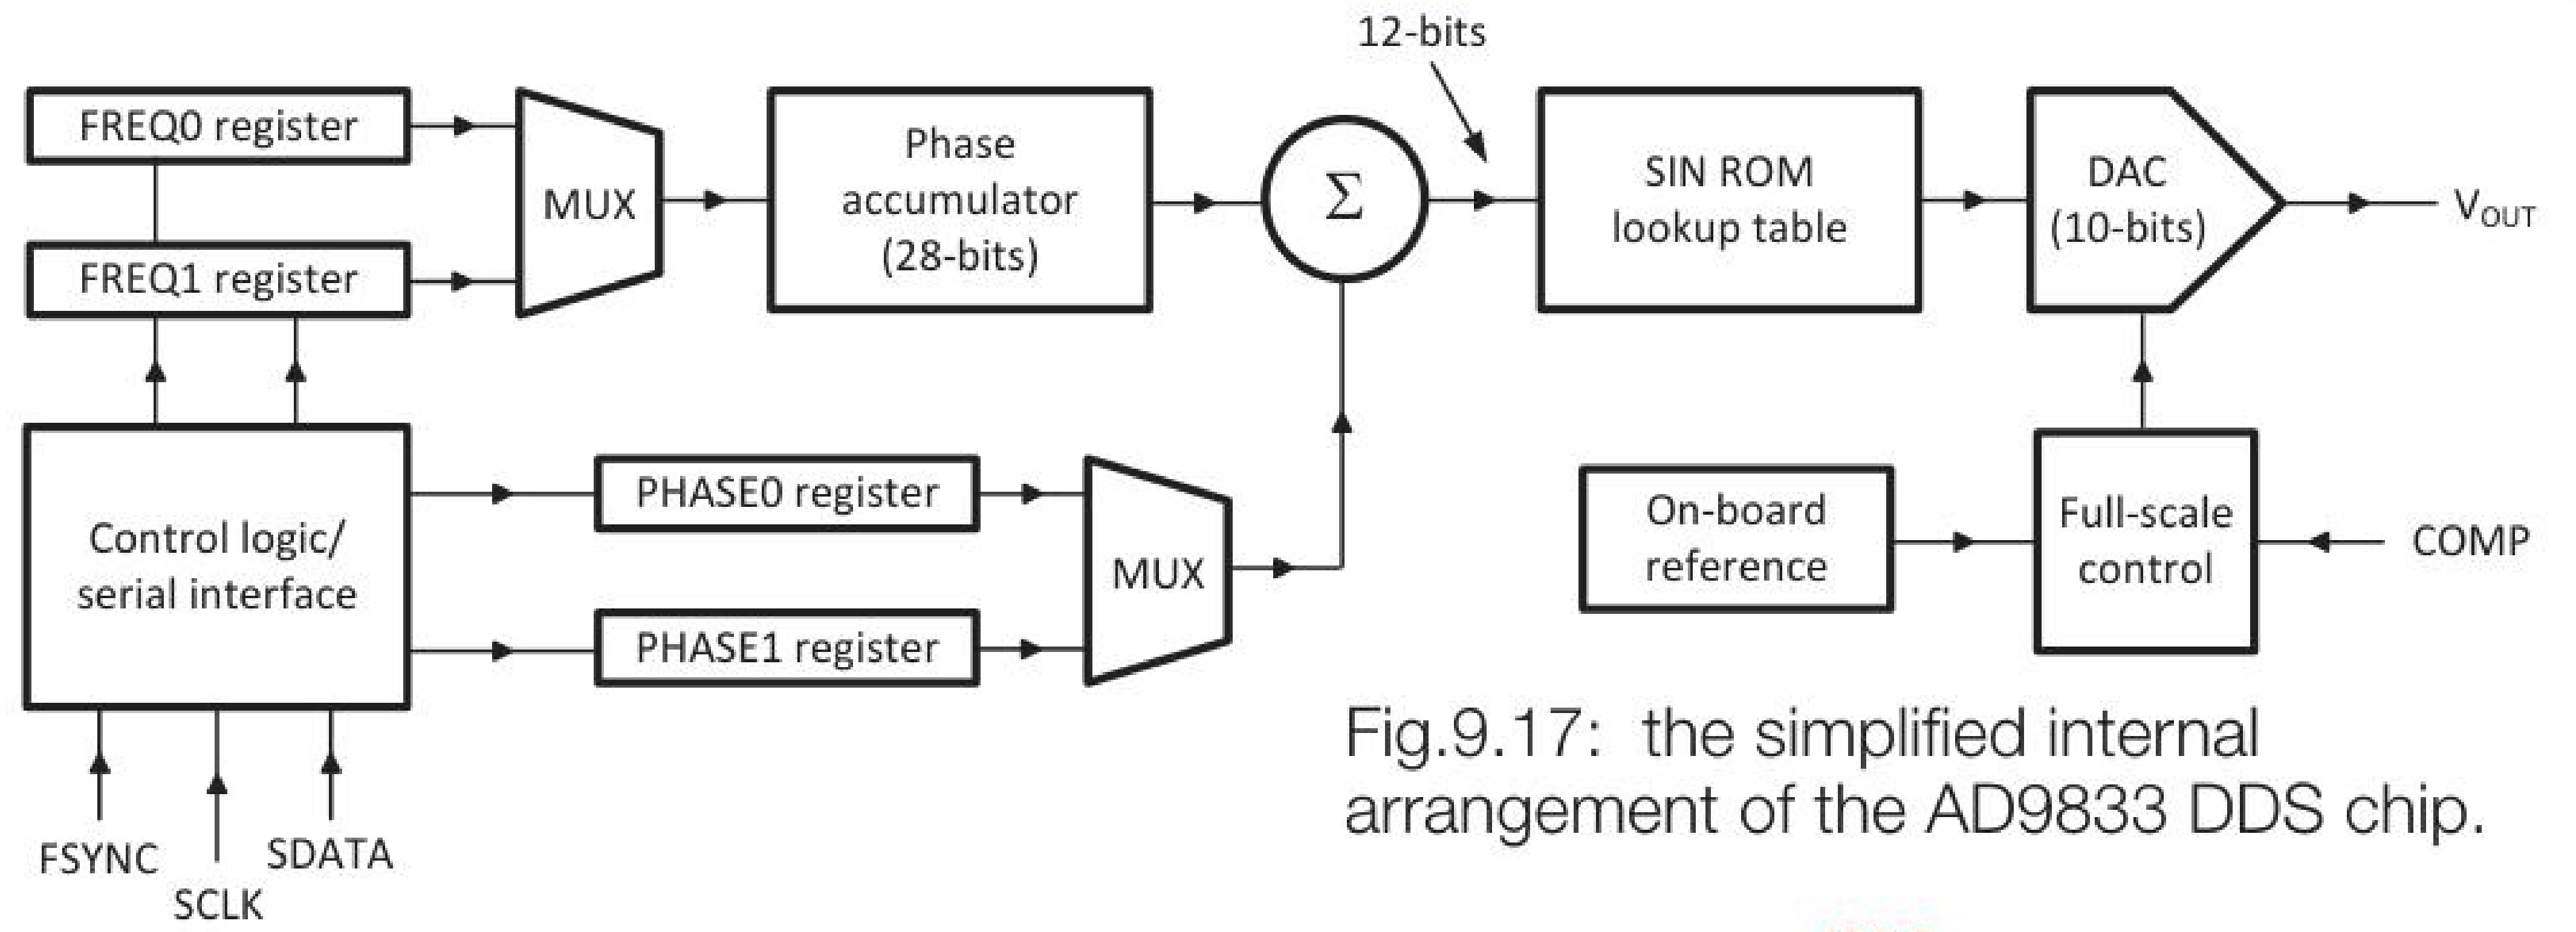
FREQ0 register
FREQ1 register
MUX
Phase accumulator (28-bits)
Σ
12-bits
SIN ROM lookup table
DAC (10-bits)
VOUT
Control logic/ serial interface
PHASE0 register
PHASE1 register
MUX
On-board reference
Full-scale control
COMP
FSYNC
SCLK
SDATA

Fig.9.17: the simplified internal arrangement of the AD9833 DDS chip.

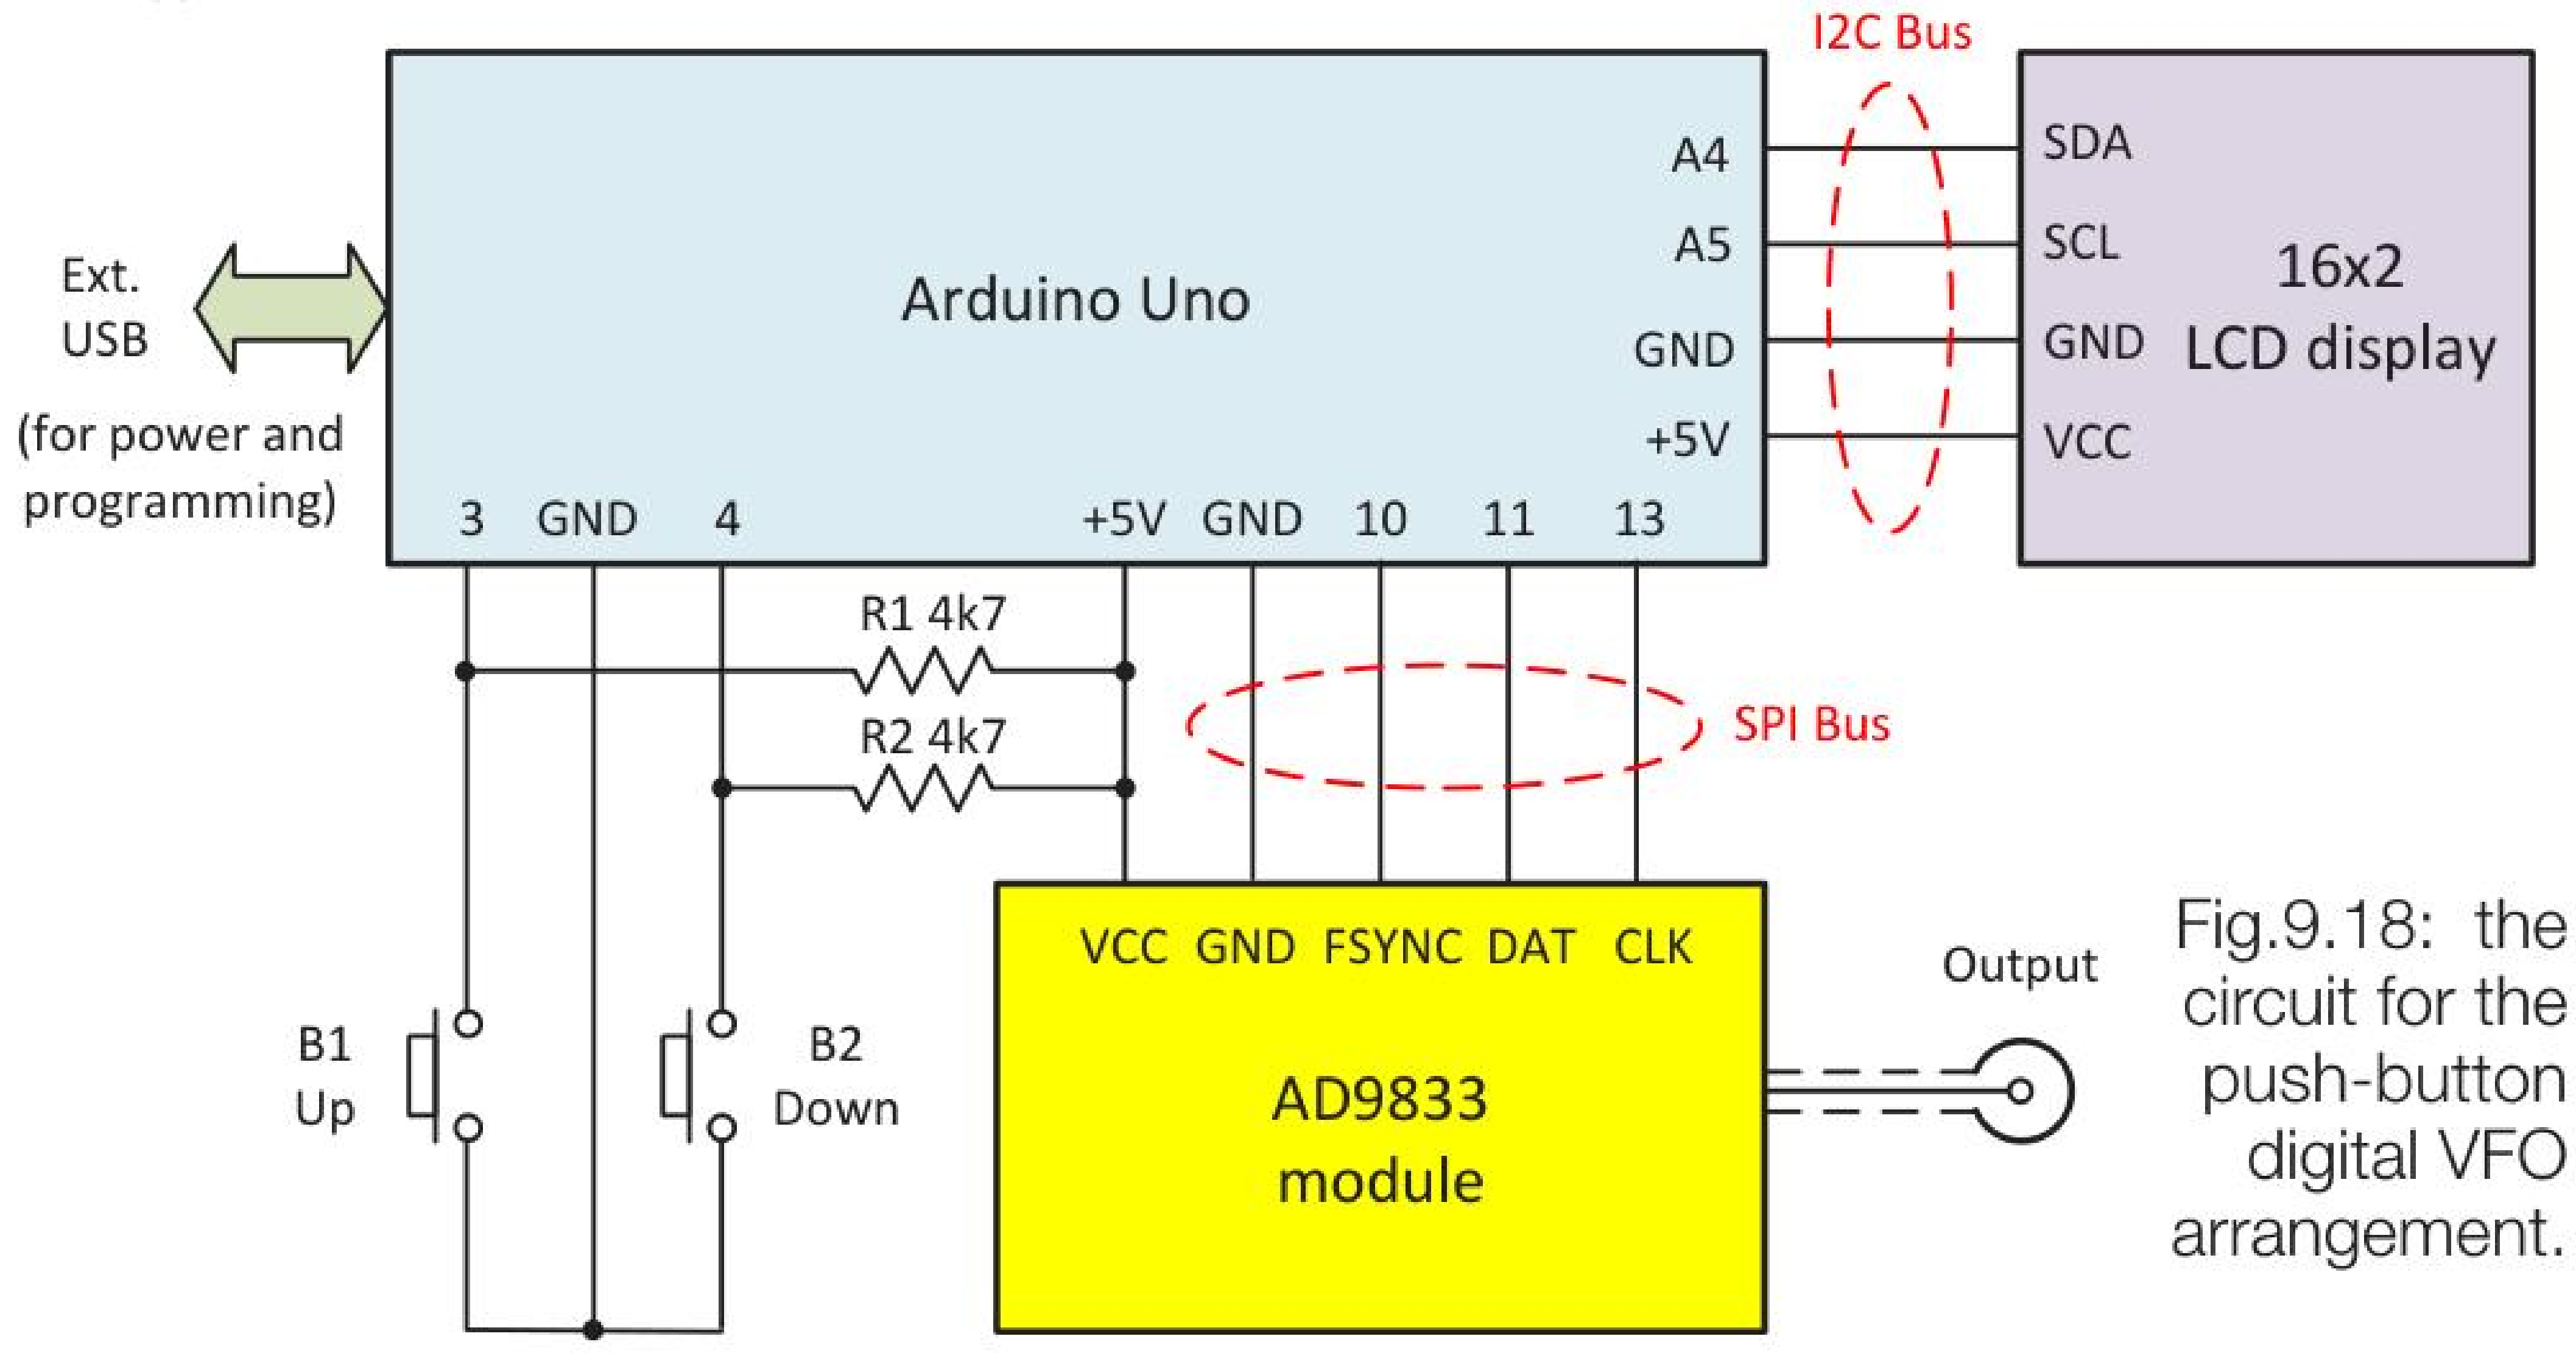
I2C Bus
Arduino Uno
A4
A5
GND
+5V
SDA
SCL
GND
VCC
16x2 LCD display
Ext. USB
(for power and programming)
3
GND
4
+5V
GND
10
11
13
R1 4k7
R2 4k7
SPI Bus
VCC GND FSYNC DAT CLK
AD9833 module
Output
B1 Up
B2 Down

Fig.9.18: the circuit for the push-button digital VFO arrangement.

more extensive example, this month's *Hands-On* project provides you with a complete Arduino Si5351 application in the form of a useful frequency marker and calibration aid.

Direct digital synthesis (DDS) and the AD9833

The AD9833 is an example of a low-cost, fully integrated DDS chip. It requires an external reference clock and a handful of other components to provide digitally synthesised sinewave outputs at up to 12.5MHz.

The frequency and phase are both software-programmable, allowing easy tuning coupled with the ability to apply digital modulation when required. Like the Si5351, the AD9833 is available in a low-cost module, such as the one shown in Fig.9.16.

The AD9833 has a standard three-wire SPI serial interface, allowing it to also be controlled by a low-cost microcontroller, although the pin that would normally be called CS (chip select) is labelled FSYNC. The pin marked CS on the module is not needed in our application.

Supplied in a 10-lead MSOP package, the chip operates from a DC supply of 2.3-5.5V. To minimise power requirements, the device incorporates a SLEEP function that disables parts of the chip that are not being used.

The simplified internal arrangement of the AD9833 is shown in Fig.9.17. It consists of two frequency select registers (FREQ0 and FREQ1), a phase accumulator, two-phase offset registers (PHASE0 and PHASE1) and a phase offset adder.

These stages are followed by a SIN ROM. This storage device uses the digital phase information as an address to a lookup table containing amplitude data.

The output of the SIN ROM is fed to a 10-bit digital-to-analog converter (DAC). This accepts successive amplitude data from the SIN ROM and converts it to corresponding analog voltages. The output from the DAC is a quantised analog waveform with an amplitude of about 0.6V peak-to-peak.

The AD9833's frequency registers are 28 bits wide and, with a 25MHz clock rate, a resolution of 0.1Hz can be achieved. For low-frequency applications with an external clock rate of 1MHz, the AD9833 can provide a resolution as low as 0.004Hz.

Although the phase accumulator uses 28 bits, its output is truncated to 12 bits because it's neither practical nor necessary to use all the bits. Using the full 28-bits would require a lookup table with a staggering 2^{28} = 268,435,456 entries!

AD9833 applications

The AD9833 can be configured for use in a wide variety of applications, not just basic RF signal or clock generation like the Si5351.

In basic frequency shift keying (FSK), the two frequency registers are loaded with different values. One represents the 'space' frequency, while the other represents the 'mark' frequency. Using the AD9833's control register, the user can then modulate the carrier frequency between the stored 'mark' and 'space' values.

The two-phase registers enable

the chip to perform phase shift keying (PSK). In this case, the carrier frequency is phase shifted, the phase being altered by an amount related to the bit stream input to the modulator. More complex modulation schemes, such as GMSK and QPSK, can also be implemented using the AD9833.

Fig.9.18 shows how an AD9833 module can be used to form the basis of a simple push-button digital VFO controlled by a few lines of code running on an Arduino Uno. The arrangement uses a 16×2 alphanumeric LCD to indicate the output frequency and two miniature push-buttons for ‘up’ and ‘down’ frequency control.

The author’s prototype digital VFO is illustrated in Fig.9.19, and the Arduino code is given in Listing 9.2. The library file developed by Marco Colli of MajicDesigns greatly simplifies interfacing with the AD9833.

The Arduino code in Listing 9.2 allows the push-button digital VFO to be adjusted in 9kHz steps over the range 7.0-7.2MHz. Values for the initial frequency, upper and lower limits can be easily changed and adapted for use in your own projects.

This application uses the I^{2}C bus to control the LCD display and the SPI interface for communication with the AD9833. Fig.9.20 shows a more advanced application in which the AD9833 forms the basis of a digitally controlled RF sweep generator with touchscreen frequency control.

Hands-On: A programmable frequency marker/calibrator

This month’s *Hands-On* project is another useful item of RF test equipment: a programmable frequency

Fig.9.19: the author’s prototype push-button digital VFO.

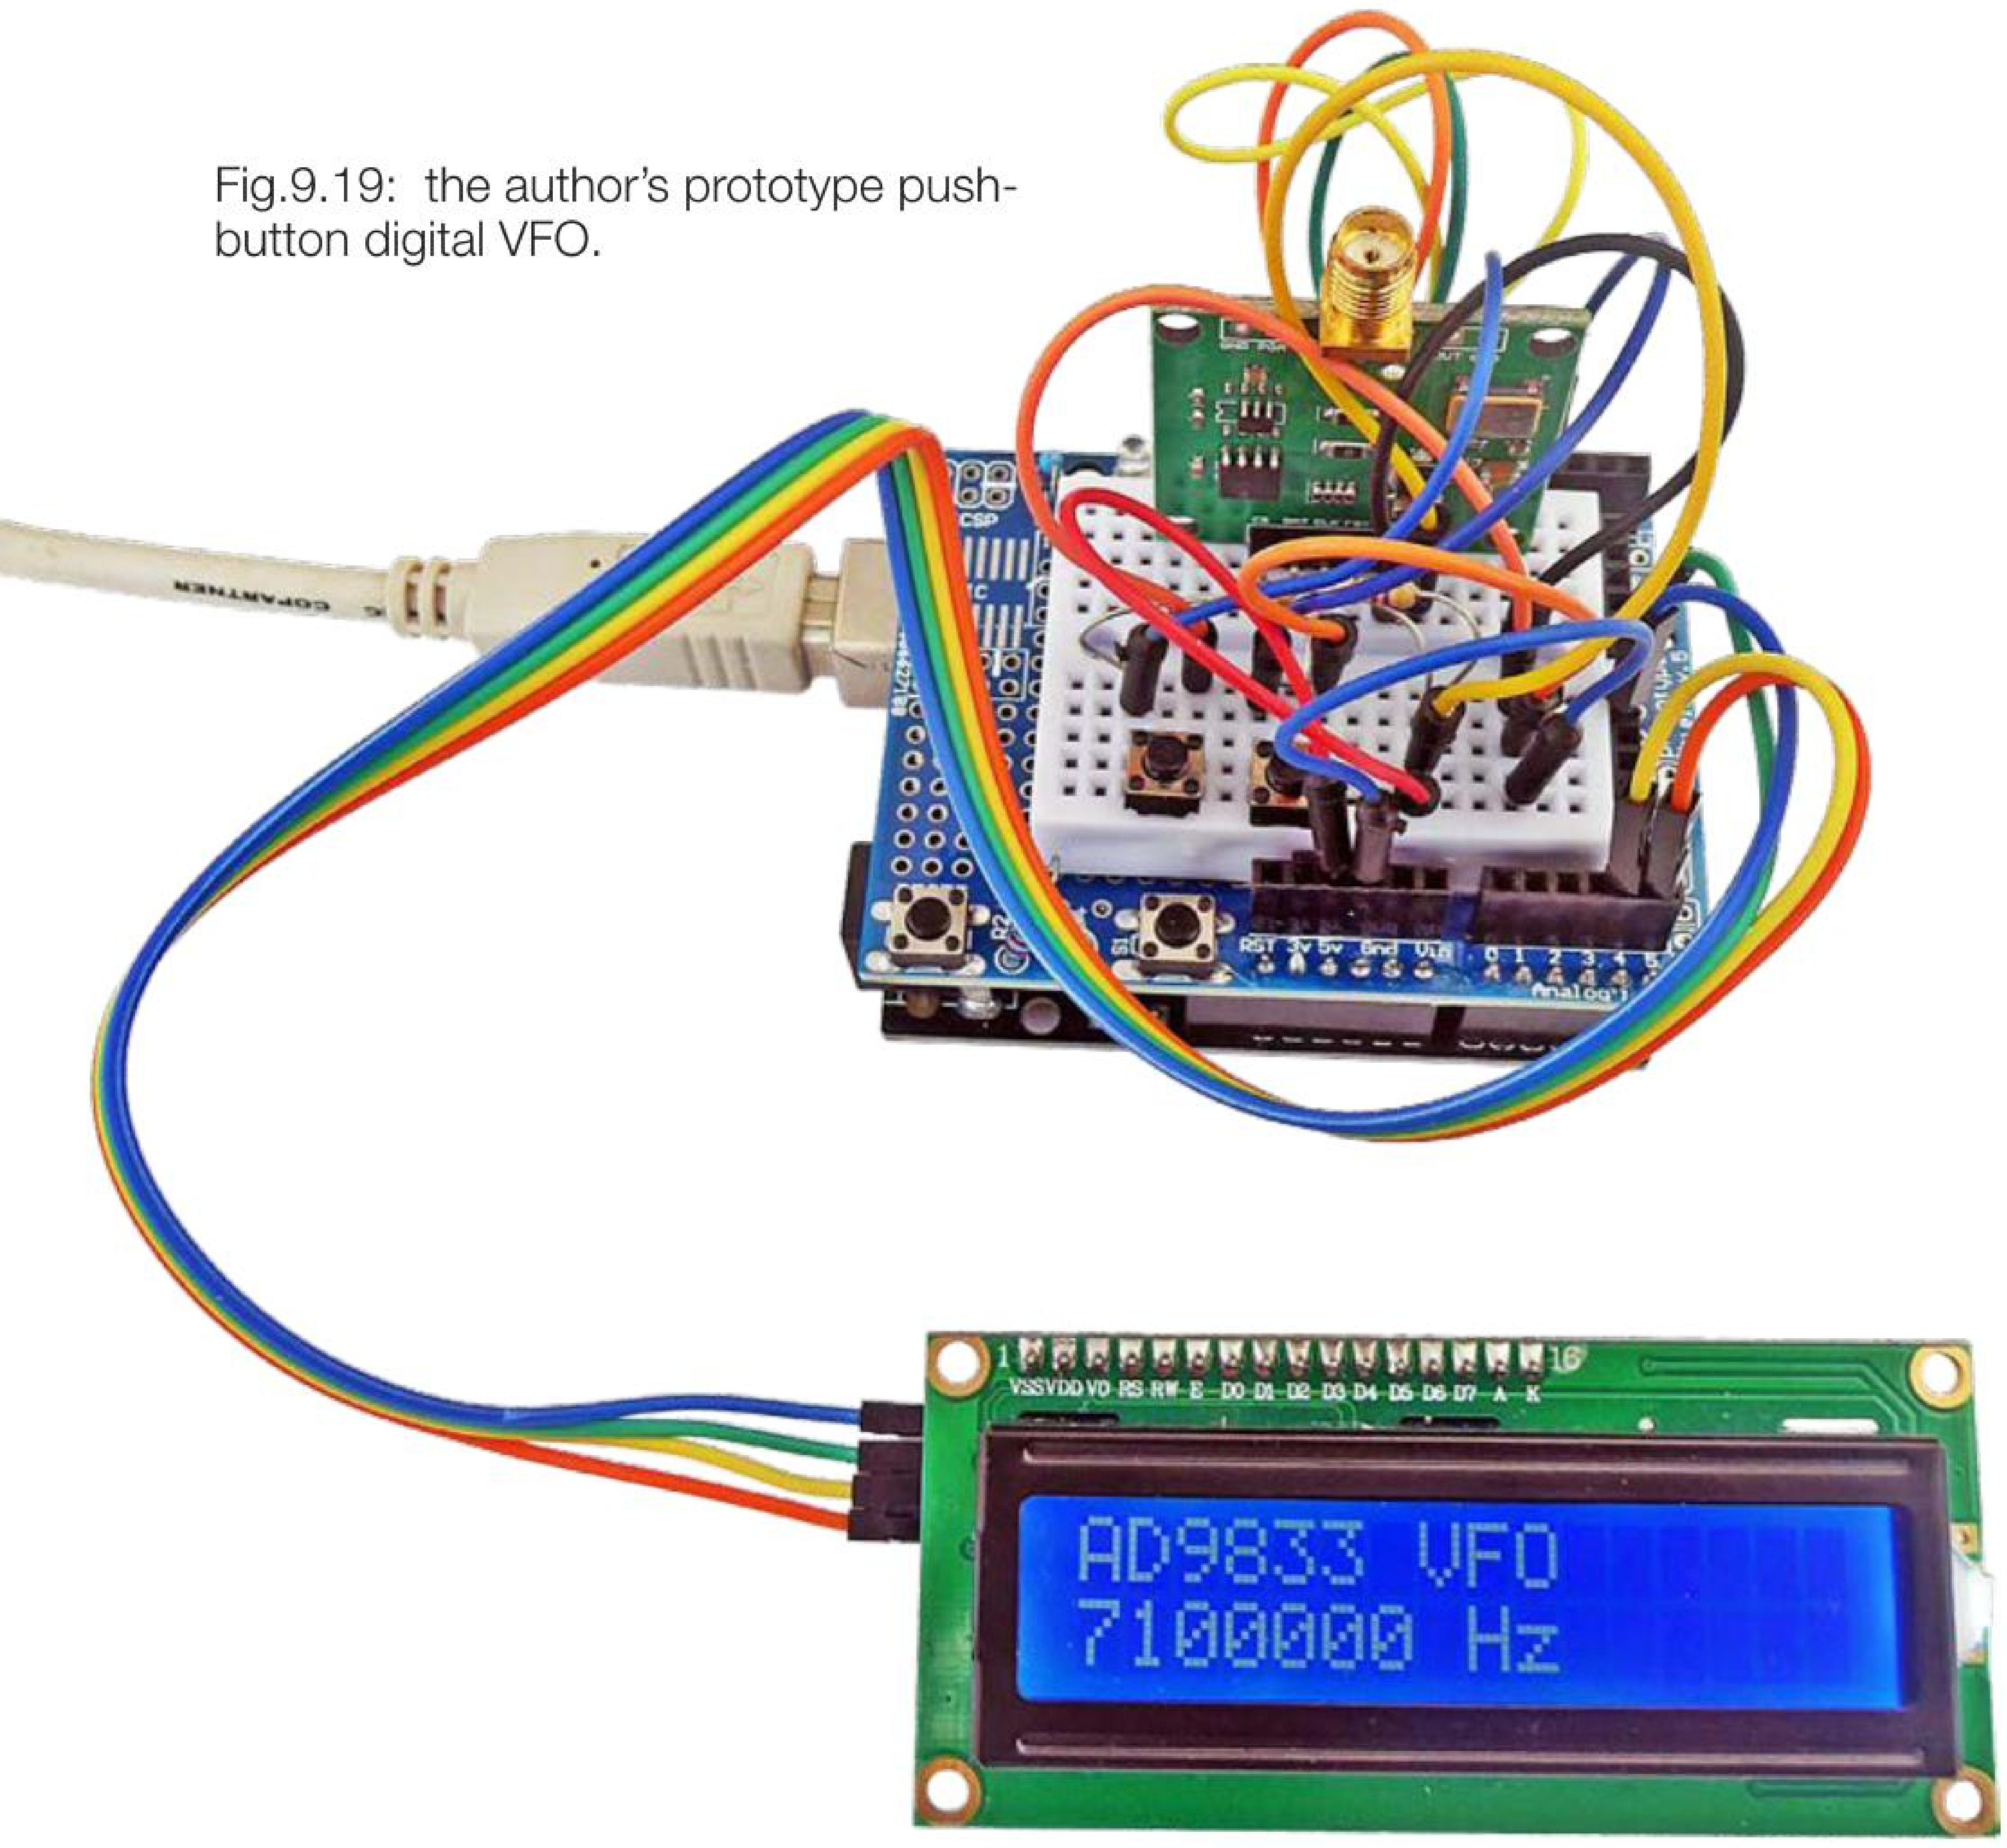

Listing 9.2: a push-button digital VFO using the AD9833.

```
#include <MD_AD9833.h>
#include <LCD-I2C.h>
#include <Wire.h>
#include <SPI.h>

unsigned long freq = 7100000;  // Initial frequency in kHz

// Step and limits
const long upperlim = 7200000;  // Upper limit in kHz
const long lowerlim = 7000000;  // Lower limit in kHz
const long step = 9000;         // Step in kHz
// Pins used for SPI data
const uint8_t PIN_DATA = 11;
const uint8_t PIN_CLK = 13;
const uint8_t PIN_FSYNC = 10;
// Buttons
const int UP = 3;
const int DOWN = 4;
//
MD_AD9833 AD(PIN_FSYNC);    // Hardware SPI
LCD_I2C lcd(0x27, 16, 2);  // Default I2C LCD

void setup(void) {
  AD.begin();
  Wire.begin();
  lcd.begin(&Wire);
  lcd.display();
  lcd.backlight();
}

void loop() {
  // Read button status
  if (digitalRead(UP) == LOW)
    freq = freq + step;
  if (digitalRead(DOWN) == LOW)
    freq = freq - step;
  // Check if limits are reached
  if (freq >= upperlim)
    freq = upperlim;
  if (freq <= lowerlim)
    freq = lowerlim;
  // Update the LCD
  lcd.setCursor(0, 0);
  lcd.print("AD9833 VFO");
  lcd.setCursor(0, 1);
  lcd.print(freq);
  lcd.print(" Hz");
  AD.setFrequency(MD_AD9833::CHAN_0, freq);
  delay(250);  // Button delay
}
```

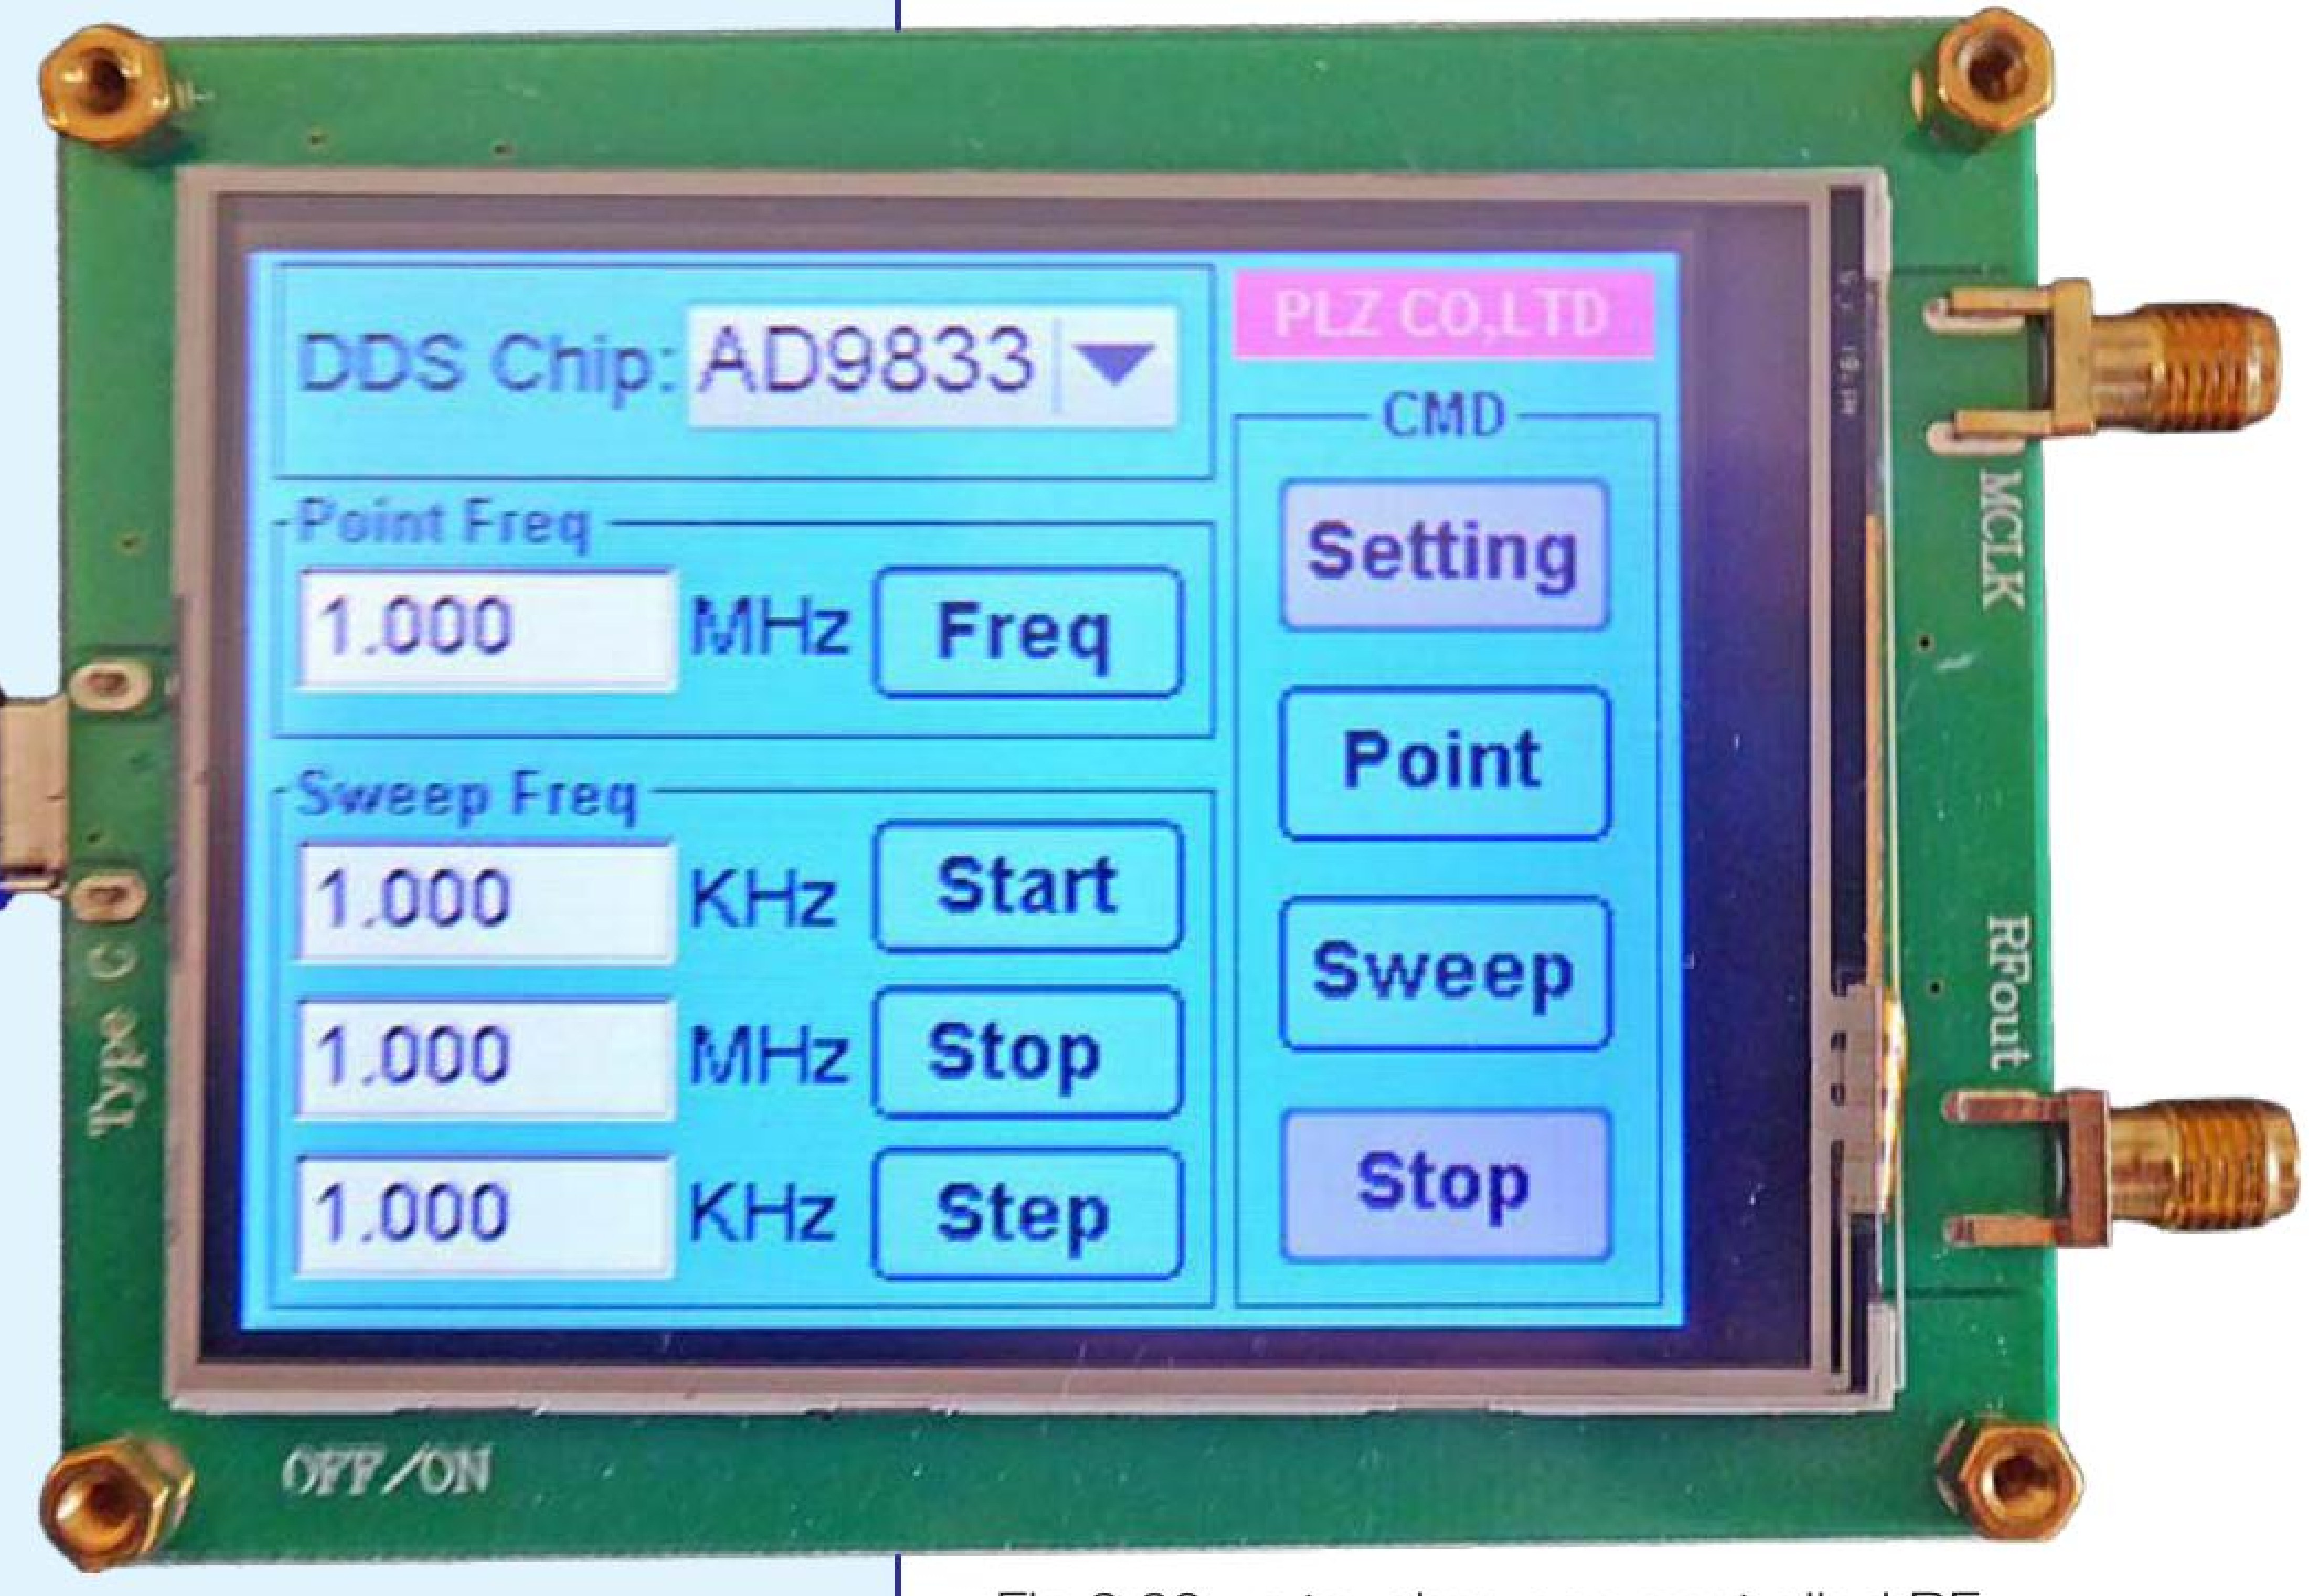

Fig.9.20: a touchscreen-controlled RF sweep generator based on the AD9833.

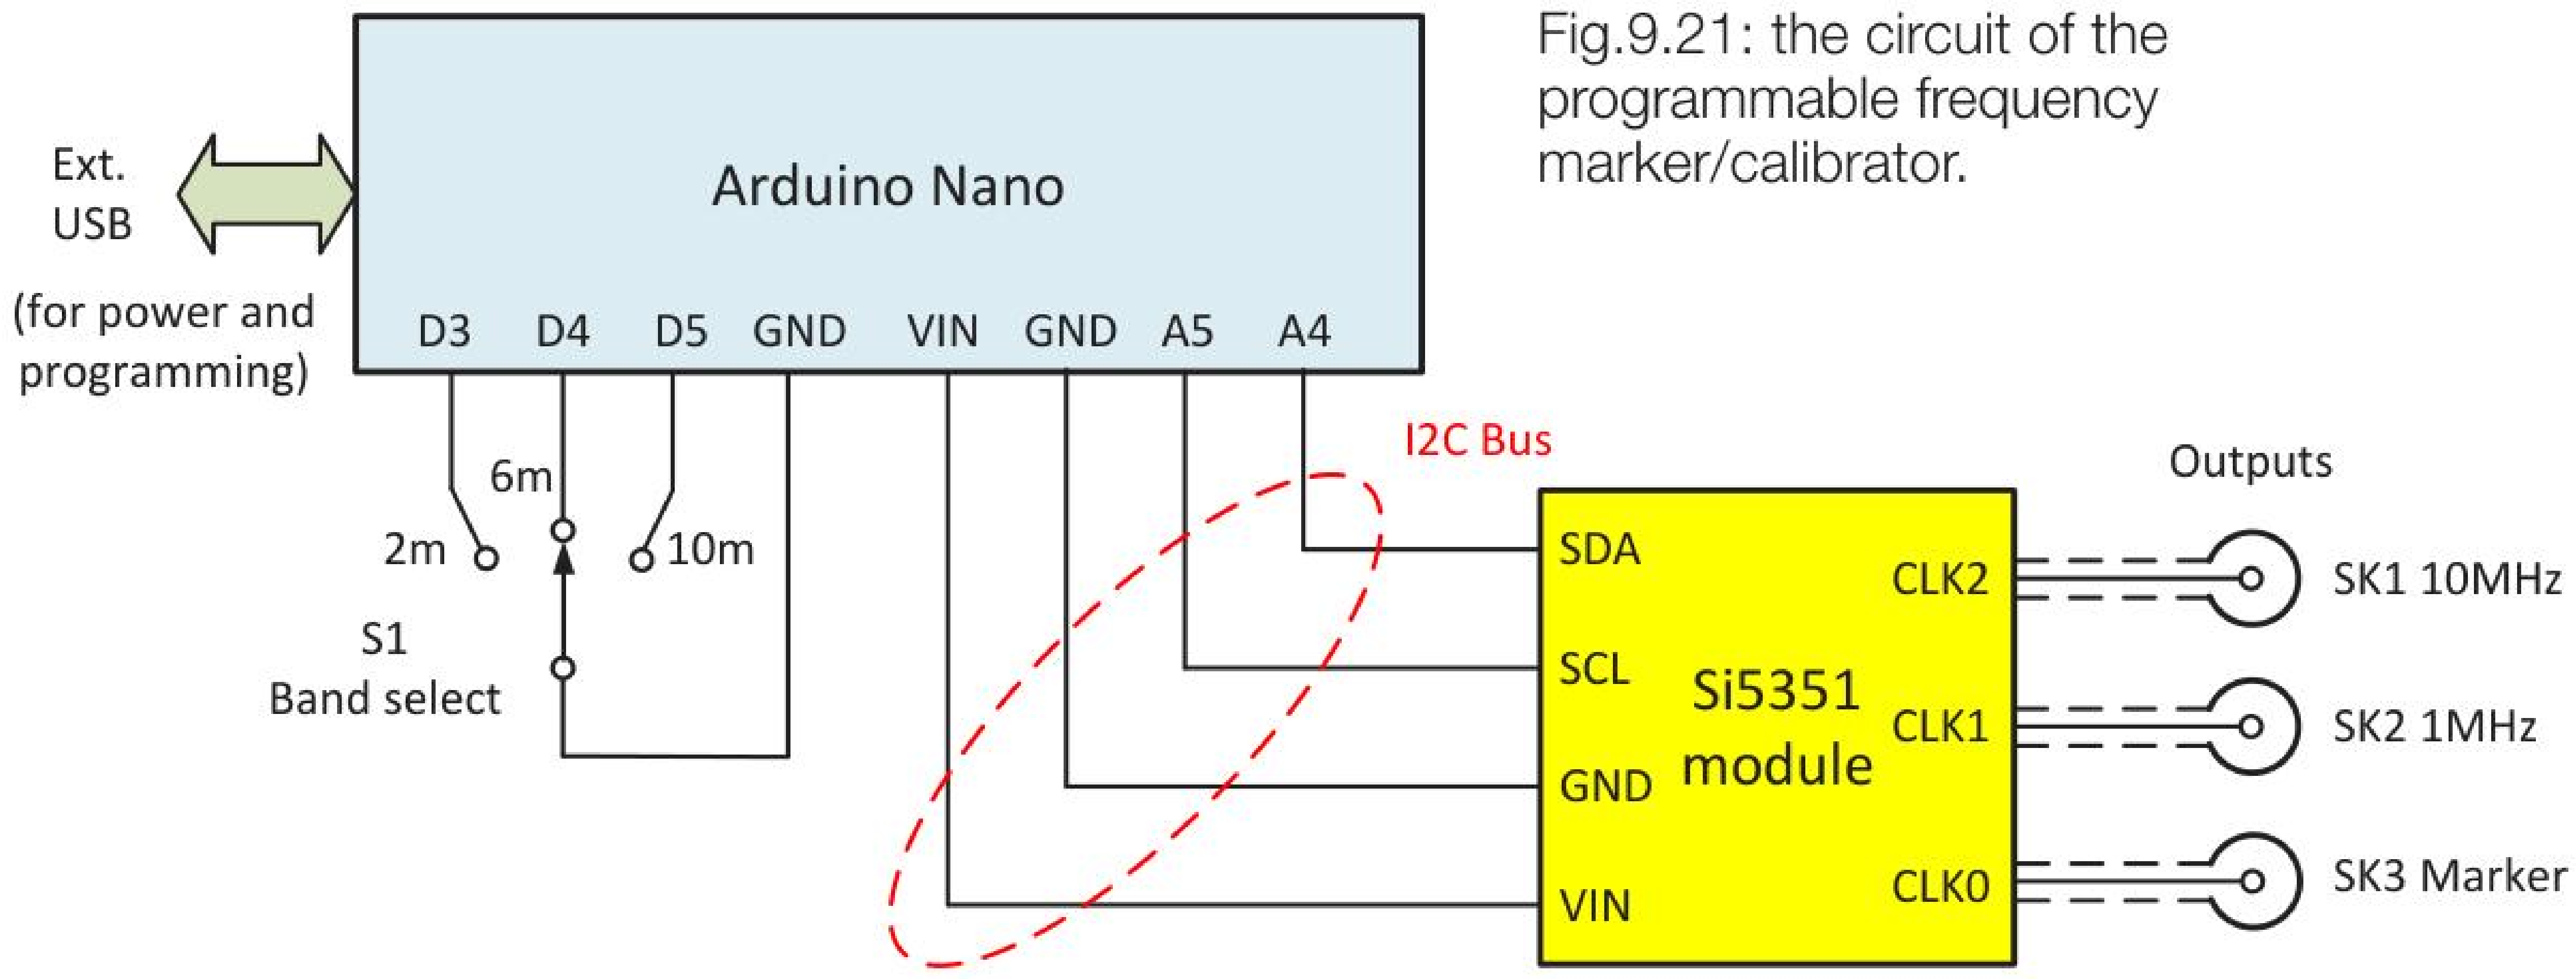
Ext. USB
(for power and programming)
Arduino Nano
D3
D4
D5
GND
VIN
GND
A5
A4
6m
2m
10m
S1
Band select
I2C Bus
SDA
SCL
GND
VIN
Si5351 module
CLK2
CLK1
CLK0
Outputs
SK1 10MHz
SK2 1MHz
SK3 Marker

Fig.9.21: the circuit of the programmable frequency marker/calibrator.

marker and calibrator, based on a low-cost Si5351 module, controlled by an Arduino Nano.

The unit provides separate calibration outputs at 1MHz and 10MHz, together with three switch-selectable marker outputs that can be programmed to frequencies of your choice. The three marker frequencies selected for the prototype were 28.050MHz in the 10m band, 50.050MHz in the 6m band and 144.500MHz in the 2m band.

The marker/calibrator circuit is shown in Fig.9.21. The Si5351 module is connected to the I^2C bus via pins A4 (SDA) and A5 (SCL). The module derives its power from the Nano's supply, VIN; the three band marker frequencies are selected by a rotary switch, S1, connected via digital inputs D3, D4 and D5.

The internal layout of the frequency marker/calibrator is shown in Fig.9.22. All components are mounted on the underside of the lid of a small ABS enclosure. To minimise wiring, the Si5351 module should be located relatively close to the Nano.

To avoid the need to fit coaxial connectors (which can be difficult unless you have the right tools), we've used ready-made panel-mounting sockets. These are available at a reasonable cost from several online suppliers and come fitted with 15cm pigtail cables and SMA connectors (see Fig.9.22).

We used standard SMA connectors for the two calibrator outputs and a single female BNC connector for the switched marker output. An alternative combination of connectors can be used to suit your requirements.

The Nano's USB supply and programming connection is made using a short 20mm USB extension lead. This mates with a bulkhead USB connector mounted on the front panel, but again, different arrangements can be made to suit your needs.

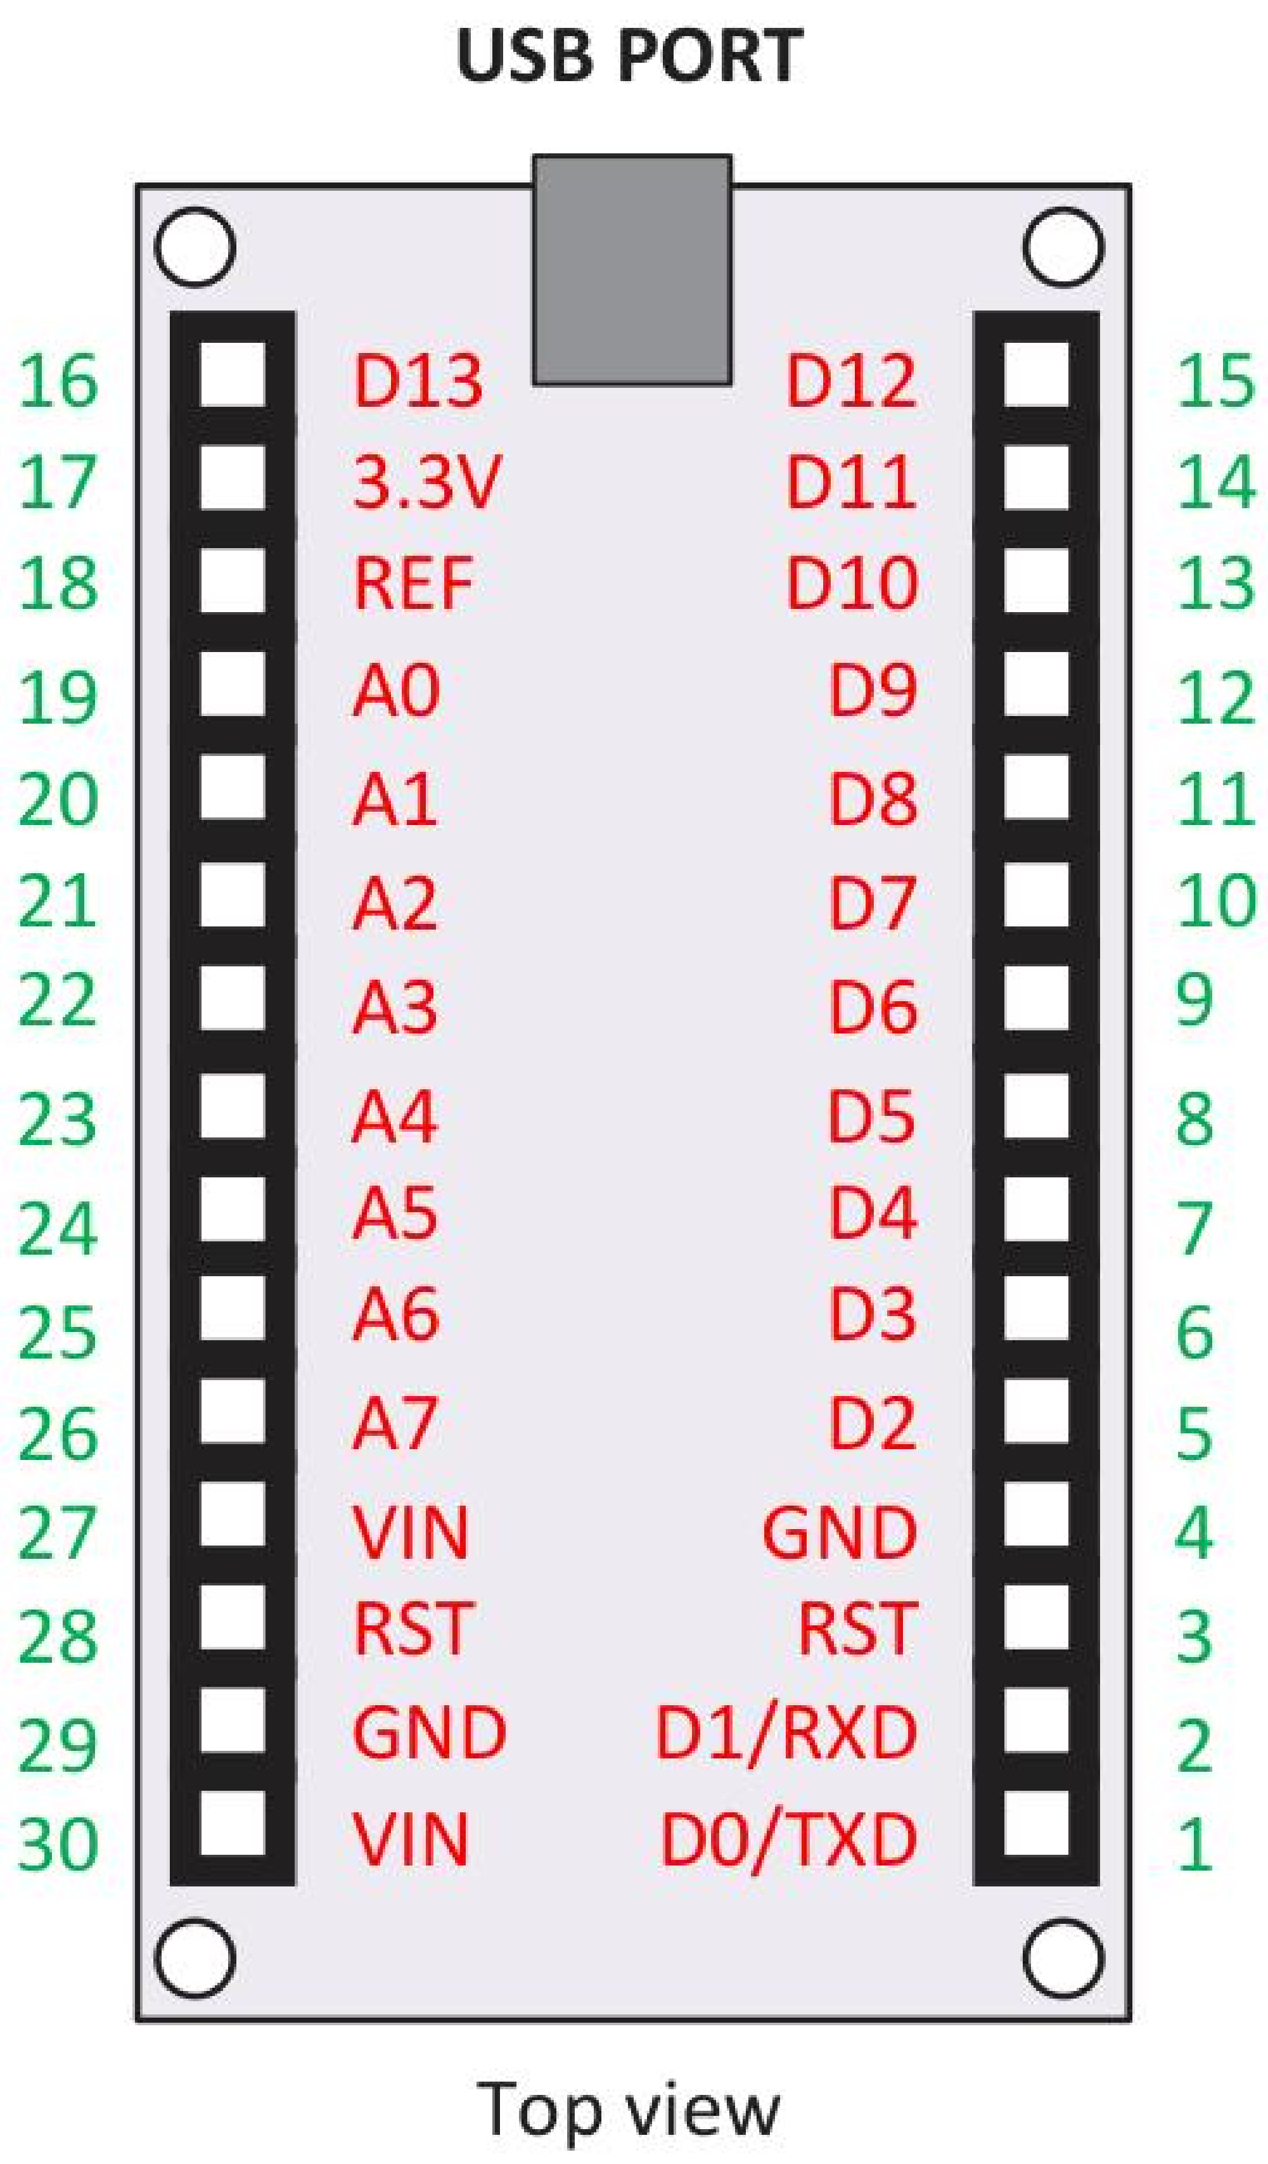
USB PORT
16
D13
D12
15
17
3.3V
D11
14
18
REF
D10
13
19
A0
D9
12
20
A1
D8
11
21
A2
D7
10
22
A3
D6
9
23
A4
D5
8
24
A5
D4
7
25
A6
D3
6
26
A7
D2
5
27
VIN
GND
4
28
RST
RST
3
29
GND
D1/RXD
2
30
VIN
D0/TXD
1
Top view

Fig.9.23: the Arduino Nano pin connections.

Wiring is straightforward using terminal pin headers fitted to both the Si5351 module and the Arduino Nano (see Fig.9.23 for the connections). The front panel layout of the frequency marker/calibrator is shown in Fig.9.24.

The Arduino code for the frequency marker/calibrator is shown in Listing 9.3. It sets the Si5351 crystal load to 6pF and the two calibrator outputs (1MHz and 10MHz)

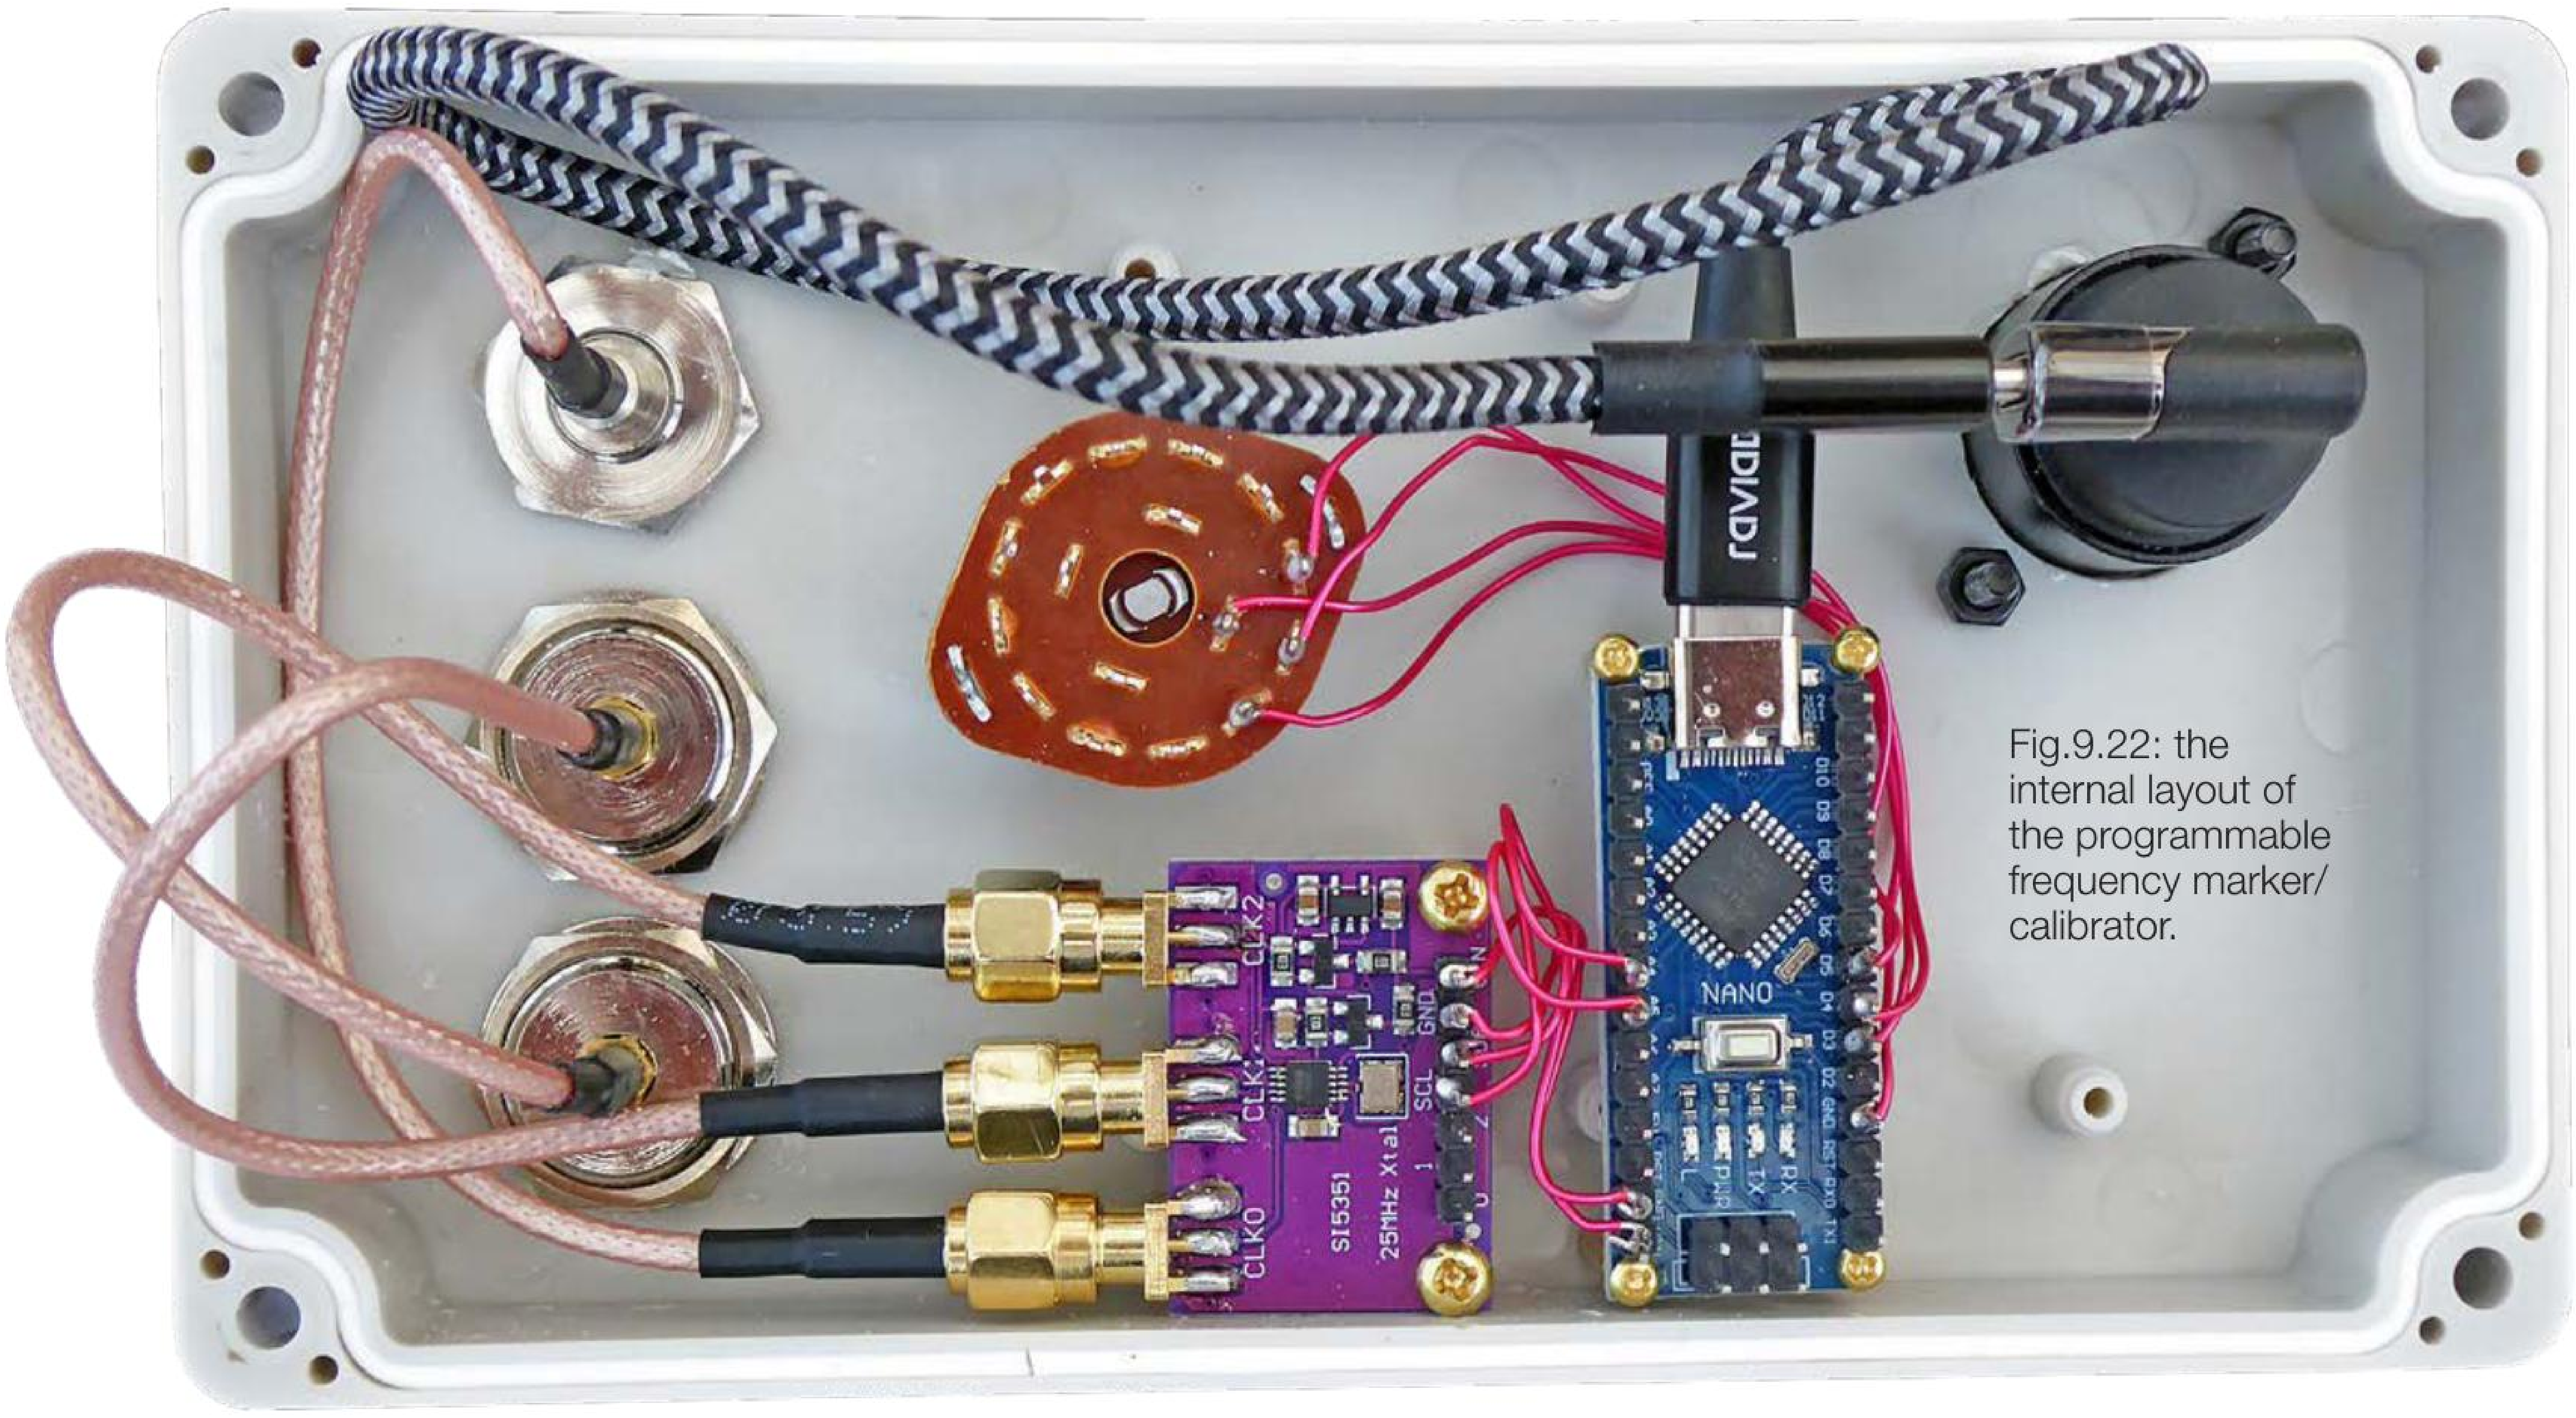
CLK2
CLK1
CLK0
Si5351
25MHz Xtal
GND
SCL
NANO

Fig.9.22: the internal layout of the programmable frequency marker/calibrator.

Parts List – programmable frequency marker/calibrator

1 small ABS enclosure, 158 × 90 × 60mm or larger
1 Arduino Nano with pin headers
1 Si5351A module
1 4-pole, 3-way rotary switch (only one pole is used) (S1)

Connectors/cables

2 panel-mounting female SMA sockets with male SMA 15cm pigtail connectors (SK1, SK2)
1 panel-mounting female BNC socket with male SMA 15cm pigtail connector (SK3)

Hardware

2 M3 × 5mm metal or nylon hexagonal standoffs/spacers
4 M2 × 5mm metal or nylon hexagonal standoffs/spacers
2 M3 × 10mm panhead machine screws
2 M3 × 10mm countersunk head machine screws
8 M2 × 10mm panhead machine screws
1 panel-mounting USB-C female socket with a male 20cm USB-C pigtail connector

to 5dBm (corresponding to about 0.4V into 50Ω). The output of the switched marker has been set to 10dBm (10mW into 50Ω) for each band, but the code can be easily modified for different output levels and frequencies.

Coming up!

In Part 10, we move on to explore short- and long-range digital wireless communication. Our *Hands-On* constructional project will be a simple RF field strength indicator. See you next time! **PE**

Listing 9.3: Arduino code for the DDS marker/calibrator.

```
#include "si5351.h"
#include "Wire.h"

Si5351 si5351;

const int SEL_144MHz = 3;  // D3 selects 2m marker
const int SEL_50MHz = 4;   // D4 selects 6m marker
const int SEL_28MHz = 5;   // D5 selects 10m marker

#define fmark1 2805000000ULL  // 10m frequency
#define fmark2 5005000000ULL  // 6m frequency
#define fmark3 14450000000ULL // 2m frequency

void setup() {
  // Enable internal pull-ups on each switch input
  pinMode(SEL_144MHz, INPUT_PULLUP);
  pinMode(SEL_50MHz, INPUT_PULLUP);
  pinMode(SEL_28MHz, INPUT_PULLUP);
  // Adjust crystal reference frequency
  si5351.init(SI5351_CRYSTAL_LOAD_6PF, 0, 0);
  // Set CLK1 to 1MHz, 5dBm
  si5351.set_freq(100000000, SI5351_CLK1);
  si5351.drive_strength(SI5351_CLK0, SI5351_DRIVE_2MA);
  // Set CLK2 to 10MHz, 5dBm
  si5351.set_freq(1000000000, SI5351_CLK2);
  si5351.drive_strength(SI5351_CLK1, SI5351_DRIVE_2MA);
  delay(500);
}

void loop() {
  // Read switch and select required marker on CLK0
  if (digitalRead(SEL_28MHz) == LOW)
    si5351.set_freq(fmark1, SI5351_CLK0);
  if (digitalRead(SEL_50MHz) == LOW)
    si5351.set_freq(fmark2, SI5351_CLK0);
  if (digitalRead(SEL_144MHz) == LOW)
    si5351.set_freq(fmark3, SI5351_CLK0);
  // Set output to approx. 10dBm
  si5351.drive_strength(SI5351_CLK0, SI5351_DRIVE_8MA);
  delay(500);
}
```

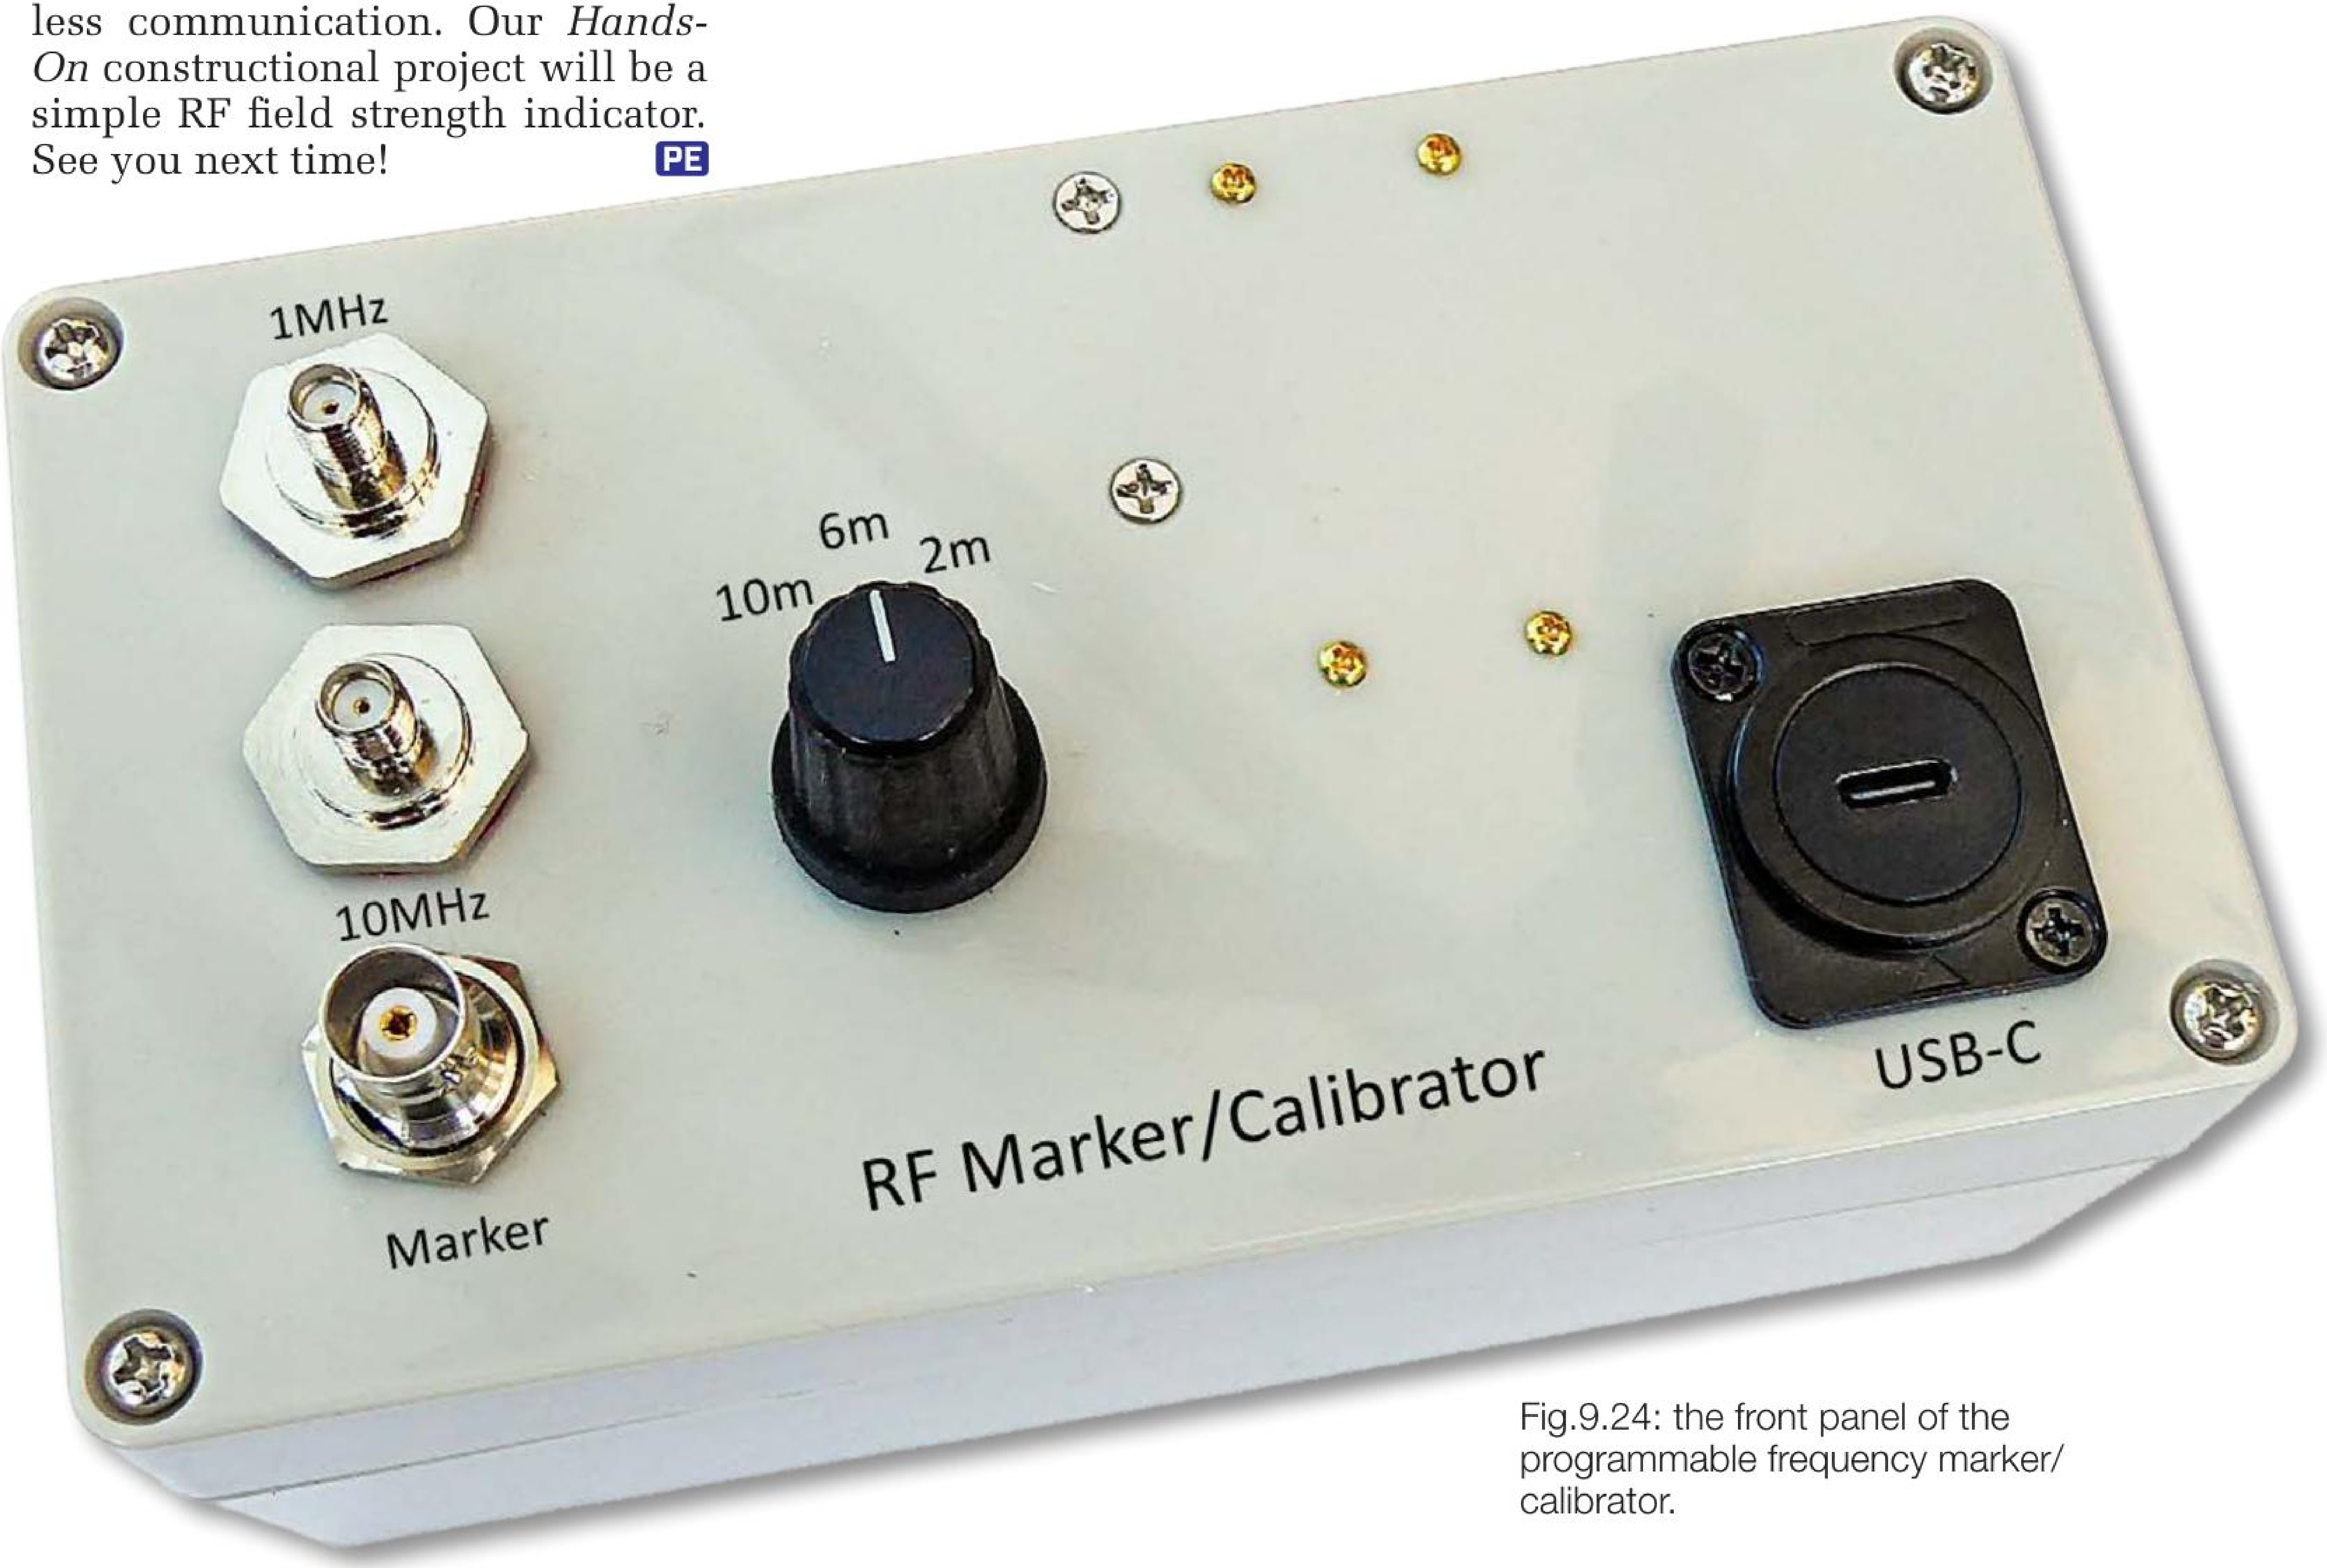

Fig.9.24: the front panel of the programmable frequency marker/calibrator.

By Steve Mansfield-Devine for PCBWay

The benefits of quality Inspection Reports for PCBs

However much you test your product designs, there will always be factors outside your control that affect your final product. One of the most important of these is the quality of the PCBs you receive from your chosen fabrication house.

All PCB manufacturers make promises about quality control and standards. It's crucial for you to know how well they live up to these assurances. That's where quality inspection reports play a vital role – they are the proof that the PCB manufacturer is meeting its promises.

Leading PCB fabs commission periodic, independent laboratory testing of their laminate materials and finished boards. These reports provide you with peace of mind and are also invaluable for ongoing process monitoring, quality assurance, regulatory compliance and certification.

PCBWay has published 14 quality inspection reports from tests carried out by Centre Testing International (CTI), an accredited third-party testing laboratory (see the bottom of the page at **www.pcbway.com/oem/quality-control.html**).

Each examination is carried out with industry-standard methods so that customers can compare the results from fabricators worldwide. This also makes them directly applicable to compliance demonstrations for IPC acceptance standards.

As a product designer or electronics engineer, you can use these reports to ensure that the physical hardware of the end product will conform to your stack-up and material set, delivering the necessary reliability, safety and longevity.

Thermomechanical tests

For example, it's essential to know how boards will respond to heat, both during assembly and when in use.

The glass transition temperature (T_g) is arguably the most important laminate specification for engineers designing boards that will be soldered with lead-free processes or that will operate at elevated temperatures. It's the temperature at which the resin binder in the laminate stops being rigid, like glass, and changes to a viscoelastic state.

This happens over a transition range, but the test results are usually reported as the midpoint. The higher the T_g temperature, the better, because exceeding it can mean the board loses dimensional stability and becomes susceptible to delamination and mechanical damage.

The common T_g figure for standard FR4 boards industry-wide is 130-140°C. However, PCBWay significantly exceeds this, with a figure of 169.61°C.

Similarly, time to delamination (T260, T288 and T300) measures how long a laminate can withstand a specific temperature before the layers separate. It is a better predictor of assembly survival than T_g alone.

The most commonly cited specification is T288, which tests the time to delamination at 288°C. The industry standard is typically 5-10 minutes for high-reliability boards, but PCBWay's tests show times consistently above the top end of the scale.

High temperatures also cause decomposition of the resin in the board, resulting in loss of mass and weakness. Most PCB fabricators aim for a decomposition temperature (T_d) of 325°C, while PCBWay achieved just over 345°C in its most recent tests.

Related to this is the coefficient of linear thermal expansion (CTE) – how much the board will expand or contract for each degree of temperature change, relative to T_g.

As the laminate is reinforced with glass cloth, expansion in the in-plane axes (X and Y) are largely constrained. However, expansion in the Z-axis (the thickness of the board) can place significant stress on plated through-holes (PTH), vias and solder joints under reflow and thermal cycling, causing barrel cracking, pad lifting and intermittent open-circuits.

For temperatures below T_g, the typical figure for CTE is 50-70ppm/°C. Again, PCBWay significantly improves on this, with a test result of 37.4ppm/°C.

These better-than-usual results mean much greater safety margins in both manufacturing and end-use for PCBWay customers.

Dimensional and structural integrity

Heat is not the only factor. There are also mechanical and structural concerns, such as board flatness (bow and twist), the integrity of inner-layer copper and interconnects, which can be checked by microsection inspection, and solder mask adhesion, which is verified by lattice tests.

These are physical aspects of the PCB that are often taken for granted. However, with bow and twist, for example, any board with SMDs that has distortion of more than 0.75% can experience problems such as tombstoning (components that should be flat on the board sticking up, attached only at one end) or the failure of solder joints on ball grid array (BGA) packages.

This is something that needs to be kept well under control, and PCBWay's tests show twist results of only 0.16% and bow of 0.12%; well below the threshold where problems typically start.

Dimensional tolerance is another factor that seems simple but can have subtle implications. You might expect that the board will be the exact size, with drill holes in the precise locations, as specified in your Gerber files. However, all manufacturing processes have certain tolerances, and you need to know that these are acceptable for your design.

The size report covers finished board dimensions, hole locations, pattern shrinkage or expansion and, in some cases, layer-to-layer registration. This helps confirm that the PCB will fit the mechanical envelope, align with mounting hardware, suit its enclosure and mate correctly with connectors, card edges and components.

In addition to reliability and manufacturing problems, poor sizing can also create signal integrity and soldering problems.

There can be compliance and certification implications, too. Factors such as safety spacing – for example, between high-voltage and low-voltage sections of the board – need to exceed the minimum allowed values in the final product.

Typical industry standards allow a ±0.1mm tolerance for the outline of the board, ±10% for the thickness, and ±50–75µm for drill position, finished hole size tolerances and layer registration tolerances.

Electrical and chemical performance

One key metric that may be on the minds of many engineers is how well their boards will stand up to high voltages.

The voltage resistant test (also known as the dielectric withstanding voltage, or hi-pot test) verifies that the PCB laminate (or solder mask over conductor pairs) can sustain 1000V DC at a controlled ramp rate of 100V/s for one minute without breakdown, flashover, sparkover or excessive leakage current. It is a simple pass/fail test.

This is an important test for the certification of mains-connected equipment. With medical devices, it's essential to ensure compliance with Means of Patient Protection (MOPP) and Means of Operator Protection (MOOP) requirements.

Related to this is the breakdown voltage. An AC voltage is applied across the test sample, ramped at 500V/s, until there is a catastrophic increase in leakage current or visible arc discharge. It provides an upper limit to the laminate's dielectric strength under AC excitation.

For power electronics designs, such as motors, uninterruptible power supplies, inverters and other cases where high-voltage conductors must be isolated on the PCB surface, solder mask breakdown voltage is a critical parameter. A designer can use these figures to derive the maximum allowable electric field strength under the solder mask and determine the minimum safe conductor spacing.

Stable layers

A PCB is far from simple. It is a sandwich of laminate substrates, traces and copper layers. How well those layers remain together is obviously crucial.

The bond strength is a measure of laminate-to-foil adhesion quality. Low adhesion is a risk that manifests as lifted pads during hand soldering or rework, delamination under thermal shock and trace peeling under mechanical vibration—all of which can result in product failures in the field.

The test measures the force required to maintain the peeling of a copper foil from the laminate, the result being the average of multiple tests. For standard FR-4 (fibreglass) boards, the general standard is 100-120N, but PCBWay's test samples averaged around 220N.

The bond strength degrades at high temperatures, which is another reason that having a high T_g value is important.

Moisture absorption also creates problems for manufacturing and long-term reliability. For instance, moisture trapped in the laminate can flash to steam during reflow, creating internal vapour pressure that causes popcorning (components flying off), delamination and blistering.

In RF and microwave PCBs, moisture absorption can be a primary material-selection criterion because absorbed water changes the laminate's dielectric properties. This can alter controlled impedances, detune RF structures such as filters and antennas, increase dielectric loss and affect phase stability, particularly in equipment exposed to humidity or temperature cycling in the field.

When a PCB absorbs water, its mass increases. The IPC-4101 maximum for standard FR-4 is a 0.32% increase, although PCBWay's figures are four times better, at 0.08%.

Production quality

There are also tests that relate directly to the production quality of the PCB. The first of these is porosity, an important factor for any board with a gold surface finish like ENIG (electroless nickel immersion gold), especially those that might be exposed to harsh environments.

Microscopic pinholes in the gold surface can expose the underlying nickel or copper, opening up a pathway for corrosion, possibly leading to 'black pad syndrome', causing poor solder joints and BGA interconnect failures that may not be detectable until deployment.

Sample boards are exposed to nitric acid vapour and then dipped in a

An operator applies solder paste to a PCB panel using a stencil in preparation for assembly.

Above: a multi-head automated PCB drilling machine.

Left: AOI (Automatic Optical Inspection) of assembled panelised PCBs.

reagent solution. The latter reacts with any exposed copper or nickel to produce visible corrosion spots. These are counted and grouped into three diameter categories: ≤0.05mm, 0.05-0.51mm and ≥0.51mm. The lower the number, the better. PCBWay's most recent report shows a perfect result of zero in all categories.

Finally, the cleanliness test measures the total quantity of ionisable (ionic) contaminants on the PCB surface. These are the products of flux activators (organic acids, halide activators), electroplating chemicals, etching solutions and handling contamination. They can result in electrochemical migration (ECM), also known as conductive anodic filament (CAF) formation and dendritic growth.

In the presence of moisture and a DC field, dissolved ions move under electromotive force and form conductive metallic filaments between adjacent conductors, causing leakage currents, intermittent short circuits and ultimately failure.

The board is washed with a solvent mixture, after which its conductivity is measured, the result being normalised to the mass of salt per unit board area. The IPC requirement is a figure of less than 1.0μg/cm^2. PCBWay's report shows a significantly better value of 0.19μg/cm^2. This is the result of a well-controlled cleaning process and minimal-residue flux due to high process control standards.

How to use the reports

These reports are not obscure technical documents but tools that exist to provide real-world data to back up assumptions engineers must make during the design process. Engineers can compare T_g, T_d, CTE, thickness, dielectric data and weave structure against their simulation and reliability assumptions about power dissipation, impedance and via life, validating their stack-up assumptions.

Size, bow/twist and micro cross-section dimensions confirm that the fabricator can meet the mechanical tolerances necessary for connectors, enclosures and heavy components. The reports also help engineers select suitable laminates.

The size, micro cross-section, CTE/T_g, bond strength and cleanliness reports provide essential data when seeking to qualify a new PCB or new supplier, especially in the automotive and aerospace sectors. The data can also support reliability assessments by helping engineers justify assumptions about PCB-related failure risks.

Signal integrity (SI) engineers frequently make use of the water absorption and micro cross-section reports to be certain that the dielectric thickness doesn't vary too much from the design specification. If it did, the trace impedance may shift, causing signal reflections and data errors in high-speed buses such as PCIe or DDR4.

OEMs can share these reports between design, quality and procurement teams to maintain a documented quality history and to justify changes in materials or suppliers during regulatory audits. Evidence of such tests may be required as part of supplier qualification or periodic audits (for example, for ISO 9001 quality management documentation).

Even when the final product is in production, the reports can aid in ongoing process monitoring. Periodic cross-sections, cleanliness and bow/twist measurements are used as process control metrics and logged in inspection reports to show statistical stability and trigger any necessary corrective actions.

Regulation and certification

Arguably the most significant benefit of these reports is in regulatory compliance and certification. For example, UL 796 (the Underwriters Laboratories standard for PCBs) requires fabricators to track T_g and bond strength.

Automotive and aerospace projects need hard evidence for thermal cycling, vibration, humidity bias and high-temperature storage. PCB-level CTE, T_g, T_d, bond strength, bow/twist, voltage and cleanliness measurements are key elements in qualification reports and control plans.

Cleanliness reports are mandatory for Class II/III medical devices under ISO 13485. Ionic contamination can lead to leakage currents that interfere with sensitive bio-signals or, in extreme cases, affect patient safety.

Similarly, safety standards such as IEC 62368-1 (IT/AV), IEC 60601-1 (medical) and their UL counterparts rely on factors such as voltage resistance/breakdown, dimensional accuracy and more to ensure that the PCB portion of the design is robust.

The micro cross-section report is the primary evidence that a board meets Class 3 requirements (such as minimum copper wrap-around at the knee of a hole) for AS9100 certification in the aerospace and defence sectors.

Conclusion

For engineers, quality inspection reports offer confirmation of technical standards. They also provide confidence that the customer's design will succeed, which is why PCBWay strives to exceed industry standards. PE

By Andrew Levido

Power Electronics

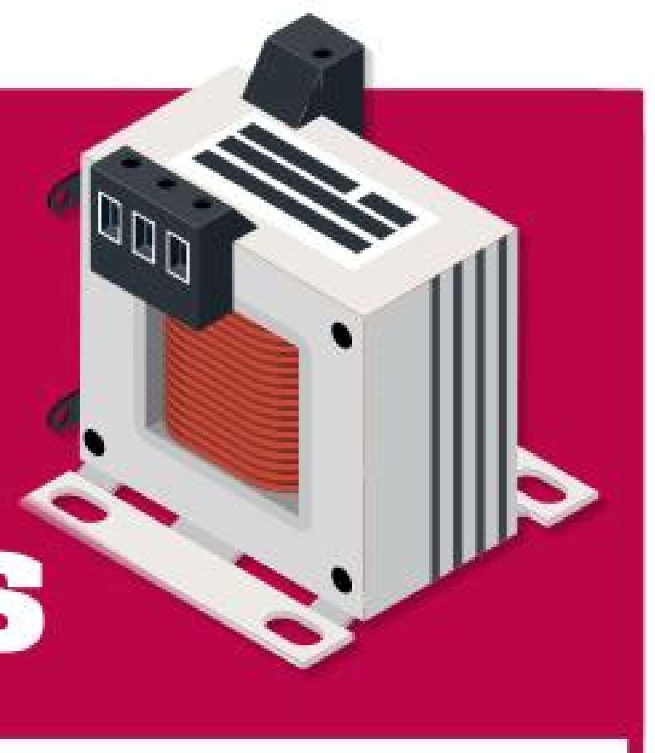

Part 4: AC-DC Conversion with Rectifiers

Power electronics as we know it today started with the invention of the mercury-arc rectifier in 1902. Rectifier-type AC-DC converters therefore predate the DC-DC converters we have been studying so far, by many decades. This month, we will look into these deceptively simple circuits.

Before the invention of the mercury-arc rectifier, the conversion of alternating current to direct current at scale required the use of rotating machines (eg, an AC motor driving a DC generator).

I think their early emergence is the reason that most power electronics textbooks and courses start with rectifiers. I have taken a different approach because I believe that rectifier circuits are more challenging to analyse than DC-DC converters.

In rectifier circuits, many of the quantities are sinusoidal or partly-sinusoidal instead of square or triangular, as they are in DC-DC converters, so calculating averages and RMS values is more difficult.

Regardless of this extra difficulty, we will take a very similar approach to analysing rectifier circuits as we did with DC-DC converters. That means we will start with the simplest possible configuration; in this case, a single-phase half-wave rectifier feeding a resistive load, as shown in Fig.1(a).

Similar to the DC-DC converter analysis, we have a source voltage and current on the left, and a load voltage and current on the right. I'll use the same conventions for AC, DC and average quantities as I have previously.

We will also use the same average value analysis technique we learned in the first article in this series. As a reminder, average value analysis is based on the fact that under periodic steady-state conditions (PSS), the average voltage across any inductor is zero and the average current through any capacitor is zero.

A single-phase half-wave rectifier

We will initially assume the voltage source and diode are ideal components. The source voltage is an AC quantity given by the expression $v_S = V_{s(pk)}\sin(\omega t)$. This $V_{s(pk)}$ term describes the amplitude of the sinusoidal voltage, and the $\sin(\omega t)$ part just describes a unit sinewave (a sinewave with an amplitude of one). The resulting voltage is shown in red on the top graph.

It is also common to describe AC quantities in terms of their 'root mean square' or RMS value. This is calculated by squaring the signal, taking the average over one cycle, and working out the square root of the result. For sinusoidal signals like our source voltage, the RMS value is just the peak value divided by $\sqrt{2}$.

The horizontal axis is also worthy of a few comments. Instead of using time as we have done so far, it is conventional to use the closely related units of phase angle (ωt, pronounced omega-t – yes, ω is the lower-case version of Ω) since trigonometric functions like sine operate on angles. If ωt is an angle and t is time, then ω must be an angular velocity, expressed in radians per second.

Angular velocity in this context is just another way of describing frequency. You can see from the graph of v_S that the length of one full cycle is 2π radians, so it should be apparent that one cycle per second (1Hz) is equal to 2π rad/s. Don't get hung up on this if it does not seem intuitive to you – it is not critical to understanding what follows.

Given that the diode can only conduct when the voltage across it is positive, the voltage seen at the load is just the positive half-cycles of the input. Since the load is resistive, the load current will have the same shape, and the source current will be the same as the load current.

The average load voltage $\langle v_L \rangle$ is harder to calculate than for the DC-DC converter where all the waveforms were square, since we have to calculate the area under the half-sine curve and divide it by the length of one cycle (2π). This requires a fairly straightforward integration, but I will spare you the details and simply give you the result, which turns out to be $\langle v_L \rangle = V_{s(pk)} \div \pi$.

Adding a filter

The output voltage is not exactly smooth DC, so to improve things, we could add an LC filter just as we did for the DC-DC converter. Before we do that, I want to add just the inductor, as in Fig.1(b), to demonstrate a point

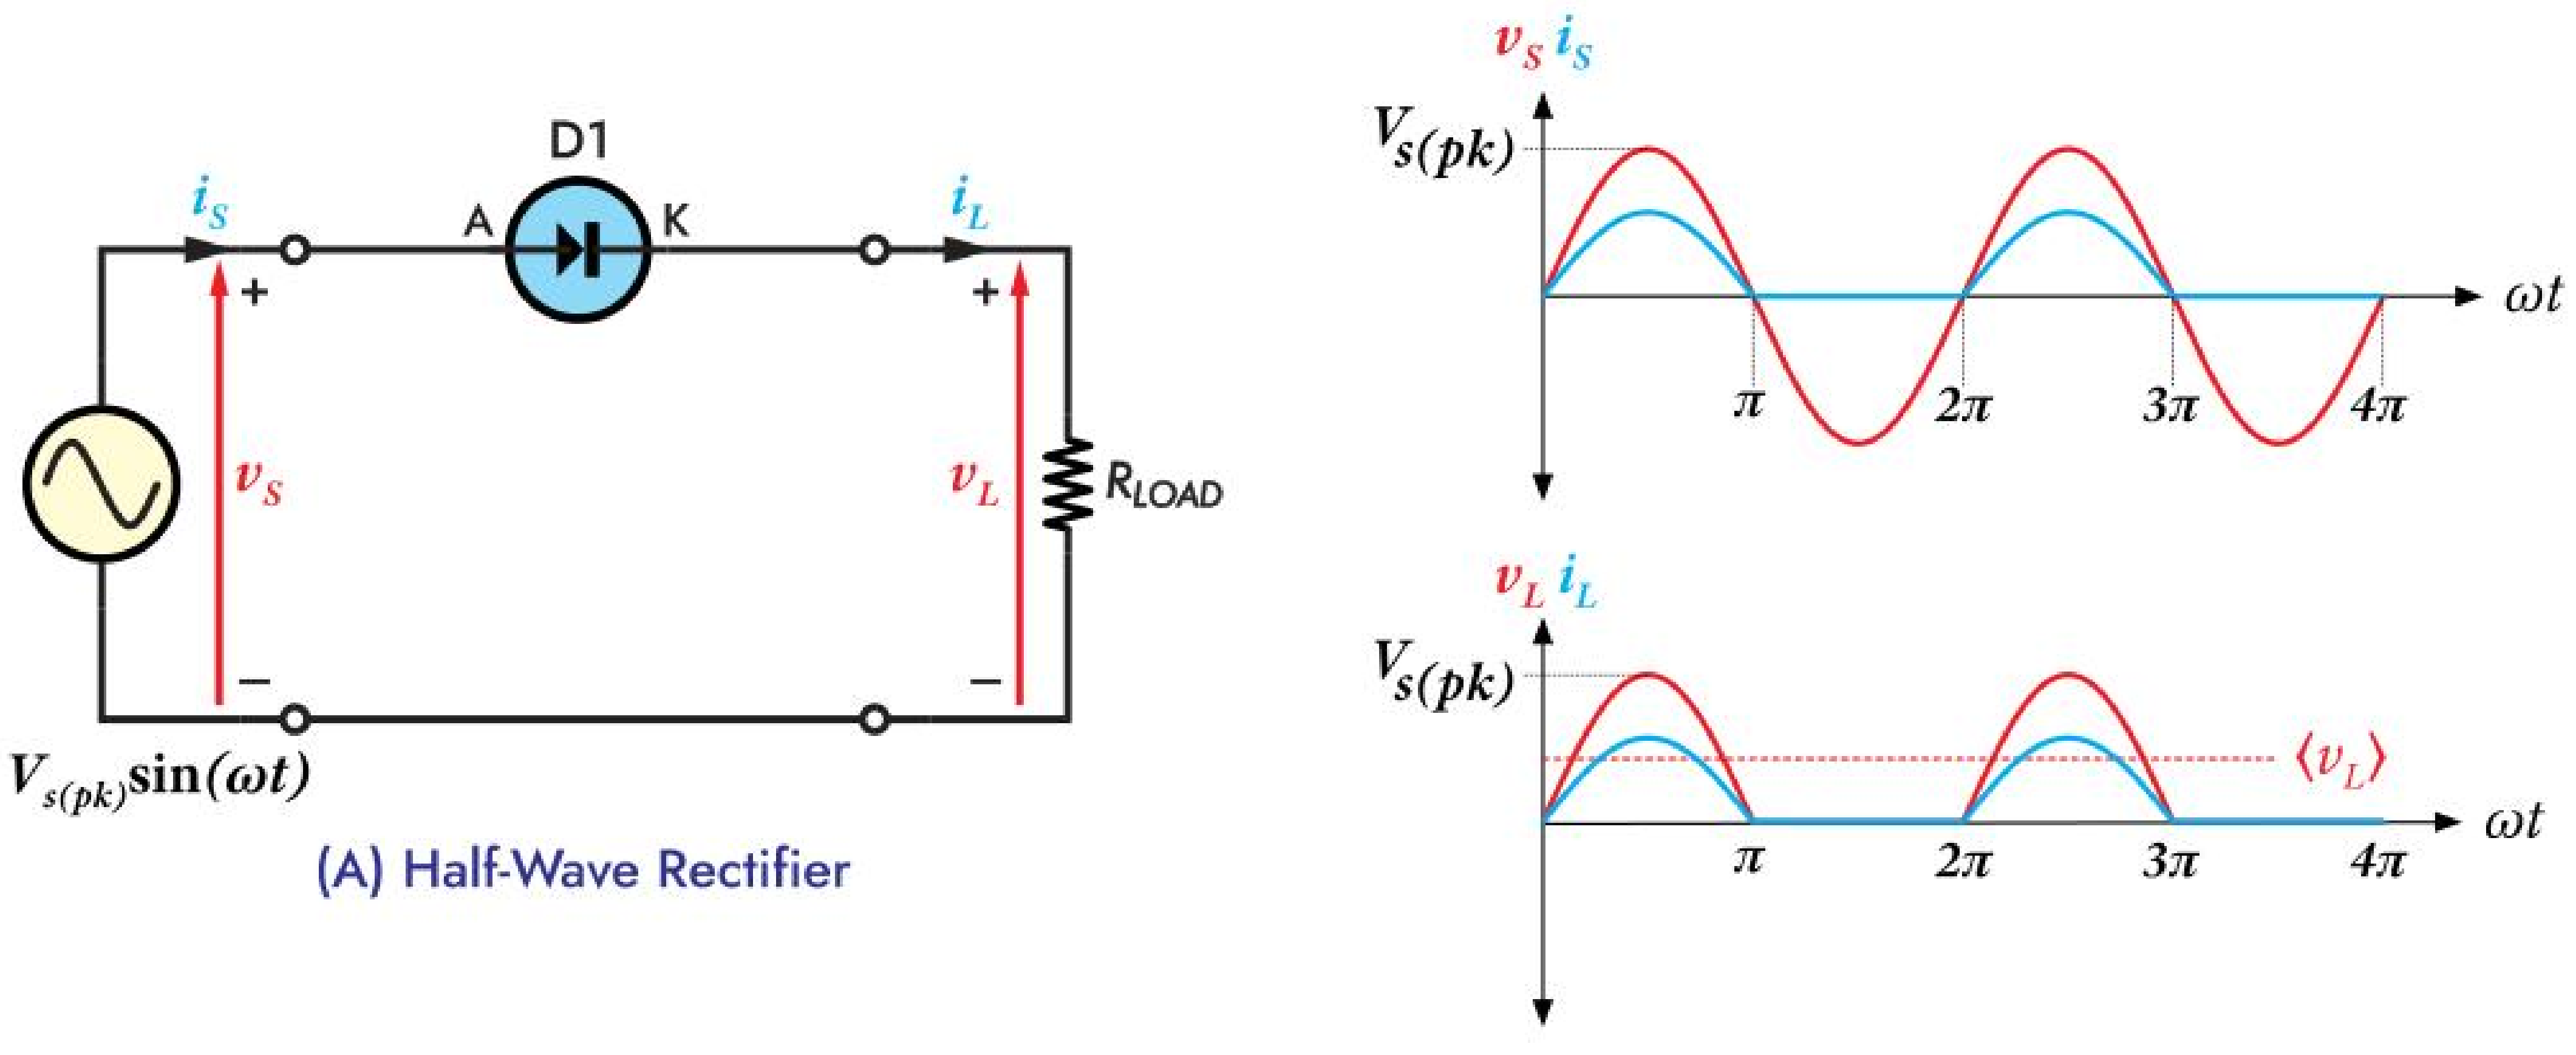

(A) Half-Wave Rectifier

Fig.1(a): half-wave rectifiers are simple but have several disadvantages. Their input current has a DC component, they have a poor power factor, and they require more filtering than their full-wave counterparts.

relating to switches like diodes that can't be switched off externally.

With the inductor in circuit, when the source positive half-cycle reaches zero at phase angle π, the inductor current is still flowing, so the diode must remain in conduction. If a current is flowing in the diode, the voltage across it must be zero. The voltage v_x must therefore go negative, following the source voltage, until the diode current falls to zero.

The load current (and source current since they are the same) is therefore 'smeared' past the end of the half-cycle.

The average voltage $\langle v_x \rangle$ will be lower than for the unfiltered rectifier because v_x is negative for a period after the zero crossing. The average load voltage $\langle v_L \rangle$ will be equal to $\langle v_x \rangle$ since the average inductor voltage must be zero due to our steady-state analysis rules.

If we were to try to increase the inductance to the point where the load (and therefore source) current is continuous, the diode would have to conduct continuously, forcing the voltage at v_x to follow the input voltage for the whole cycle. Under these circumstances, the average voltage at v_x would be zero, so the current would also be forced to zero.

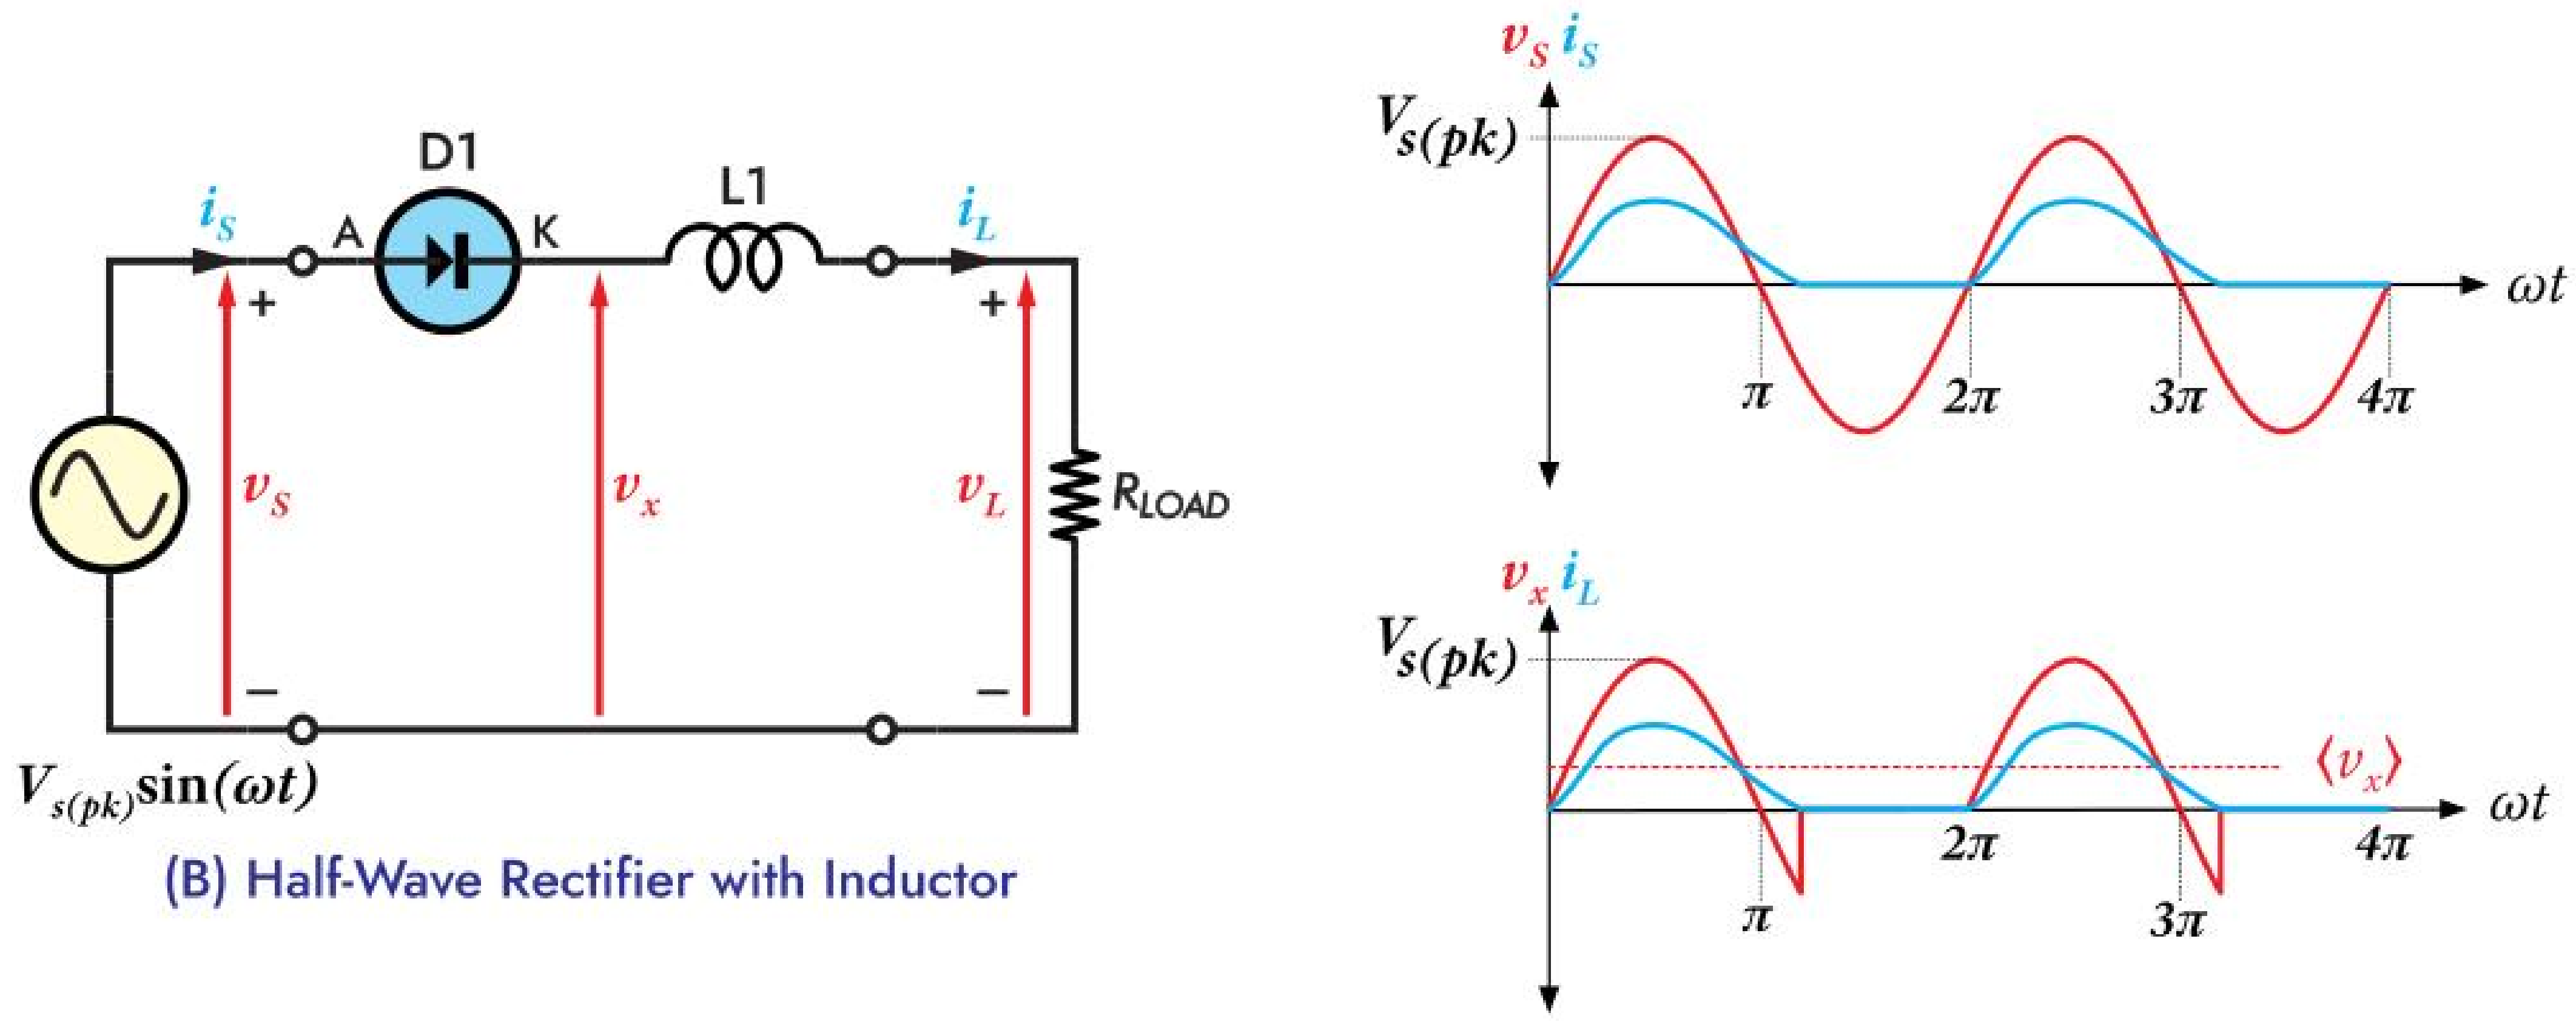
D1
L1
i_S
A
K
i_L
v_S
v_x
v_L
RLOAD
$V_{s(pk)}\sin(\omega t)$
(B) Half-Wave Rectifier with Inductor
v_S i_S
$V_{s(pk)}$
ωt
π
2π
3π
4π
v_x i_L
$\langle v_x \rangle$

Fig.1(b): a series inductor provides smoothing, but introduces new problems.

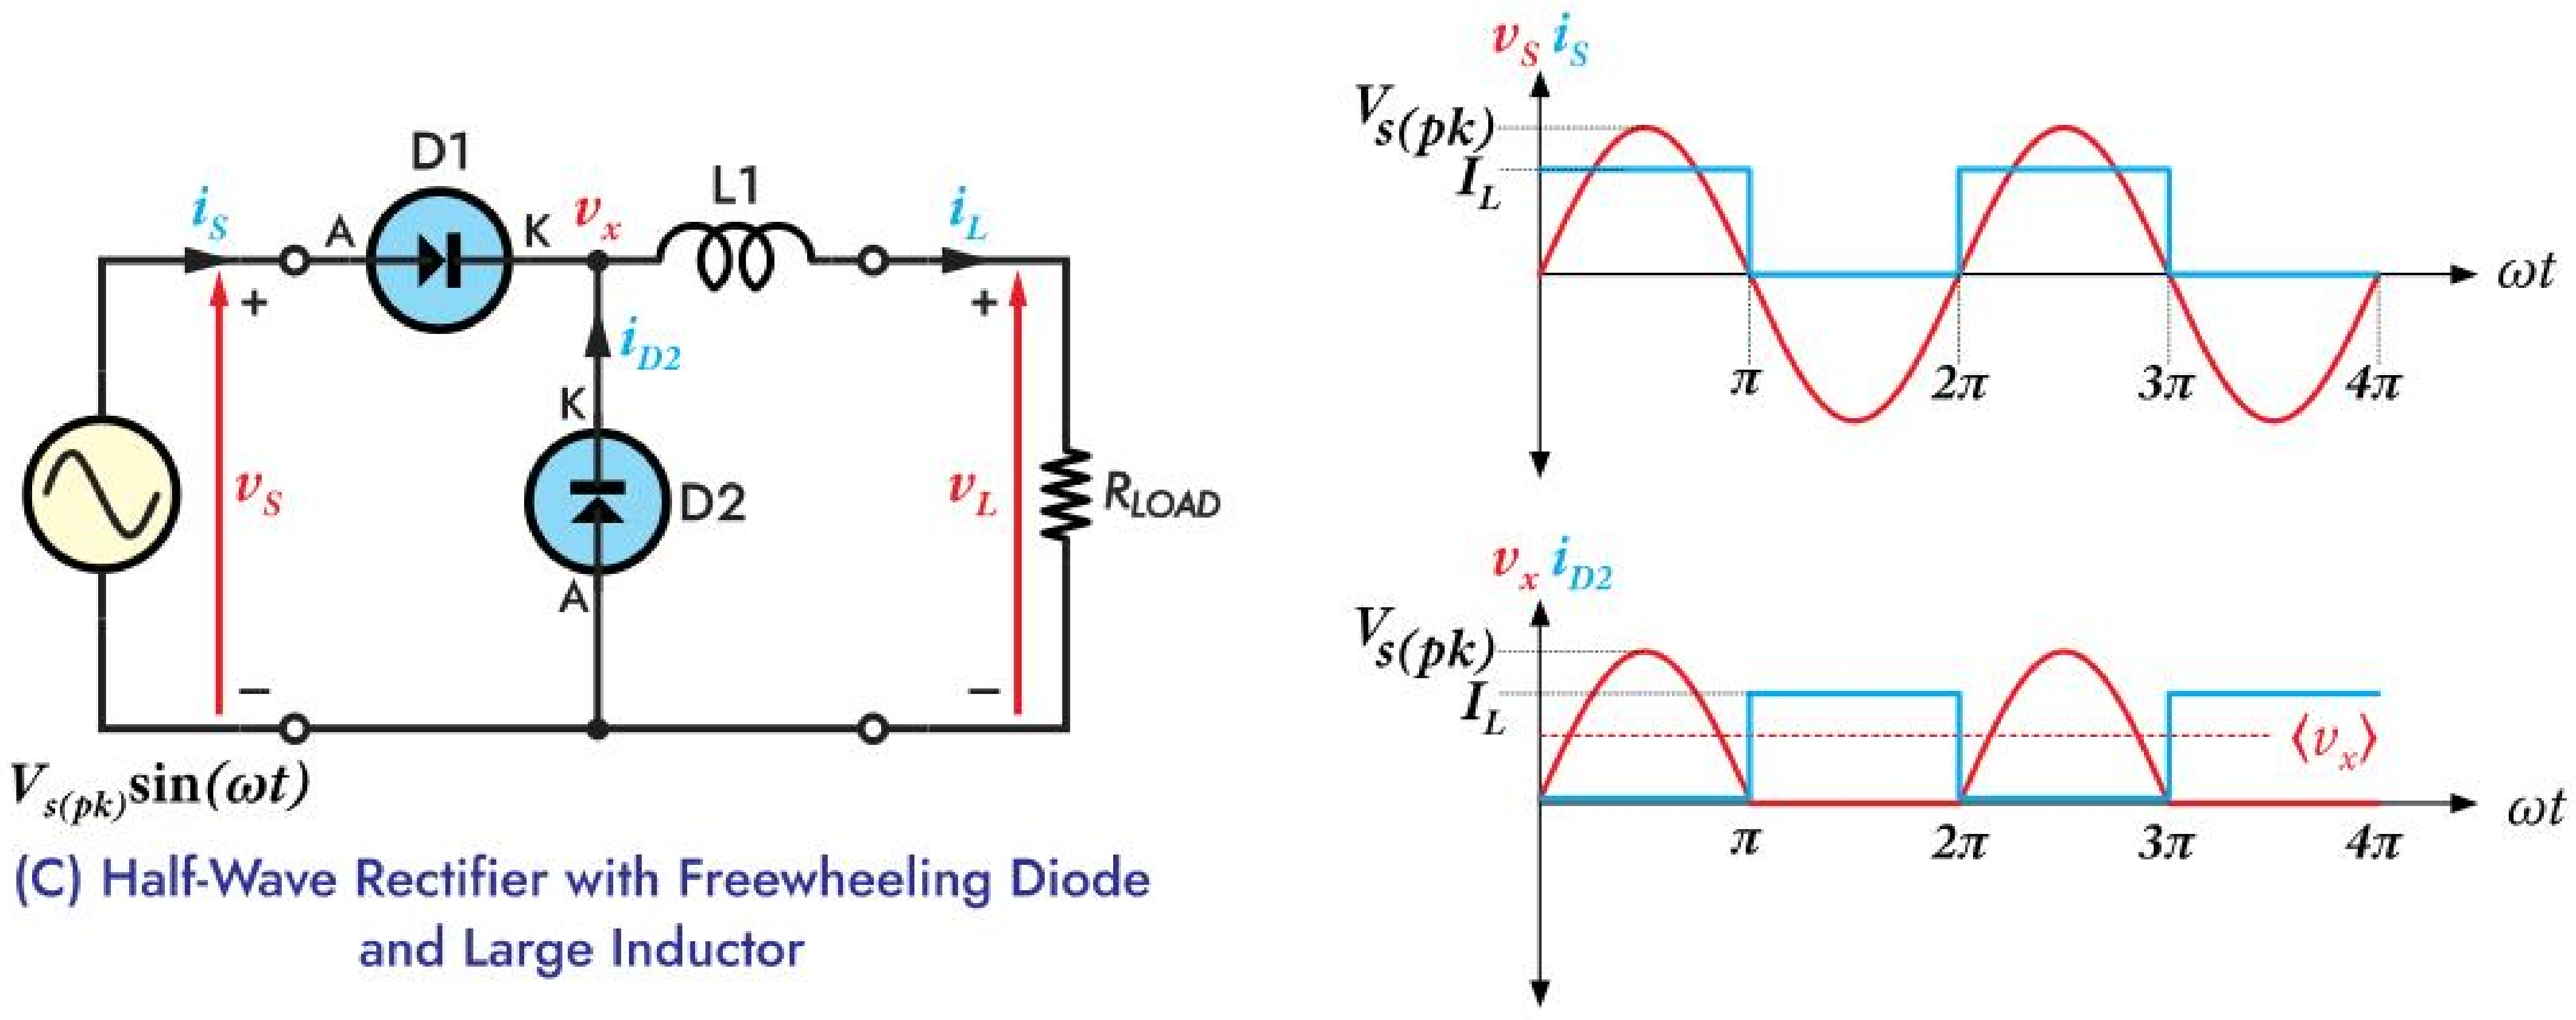
D1
L1
i_S
A
K
v_x
i_L
i_{D2}
D2
v_S
v_L
RLOAD
$V_{s(pk)}\sin(\omega t)$
(C) Half-Wave Rectifier with Freewheeling Diode and Large Inductor
v_S i_S
$V_{s(pk)}$
I_L
ωt
π
2π
3π
4π
v_x i_{D2}
$\langle v_x \rangle$

Fig.1(c): a second 'freewheeling' diode solves most of those problems.

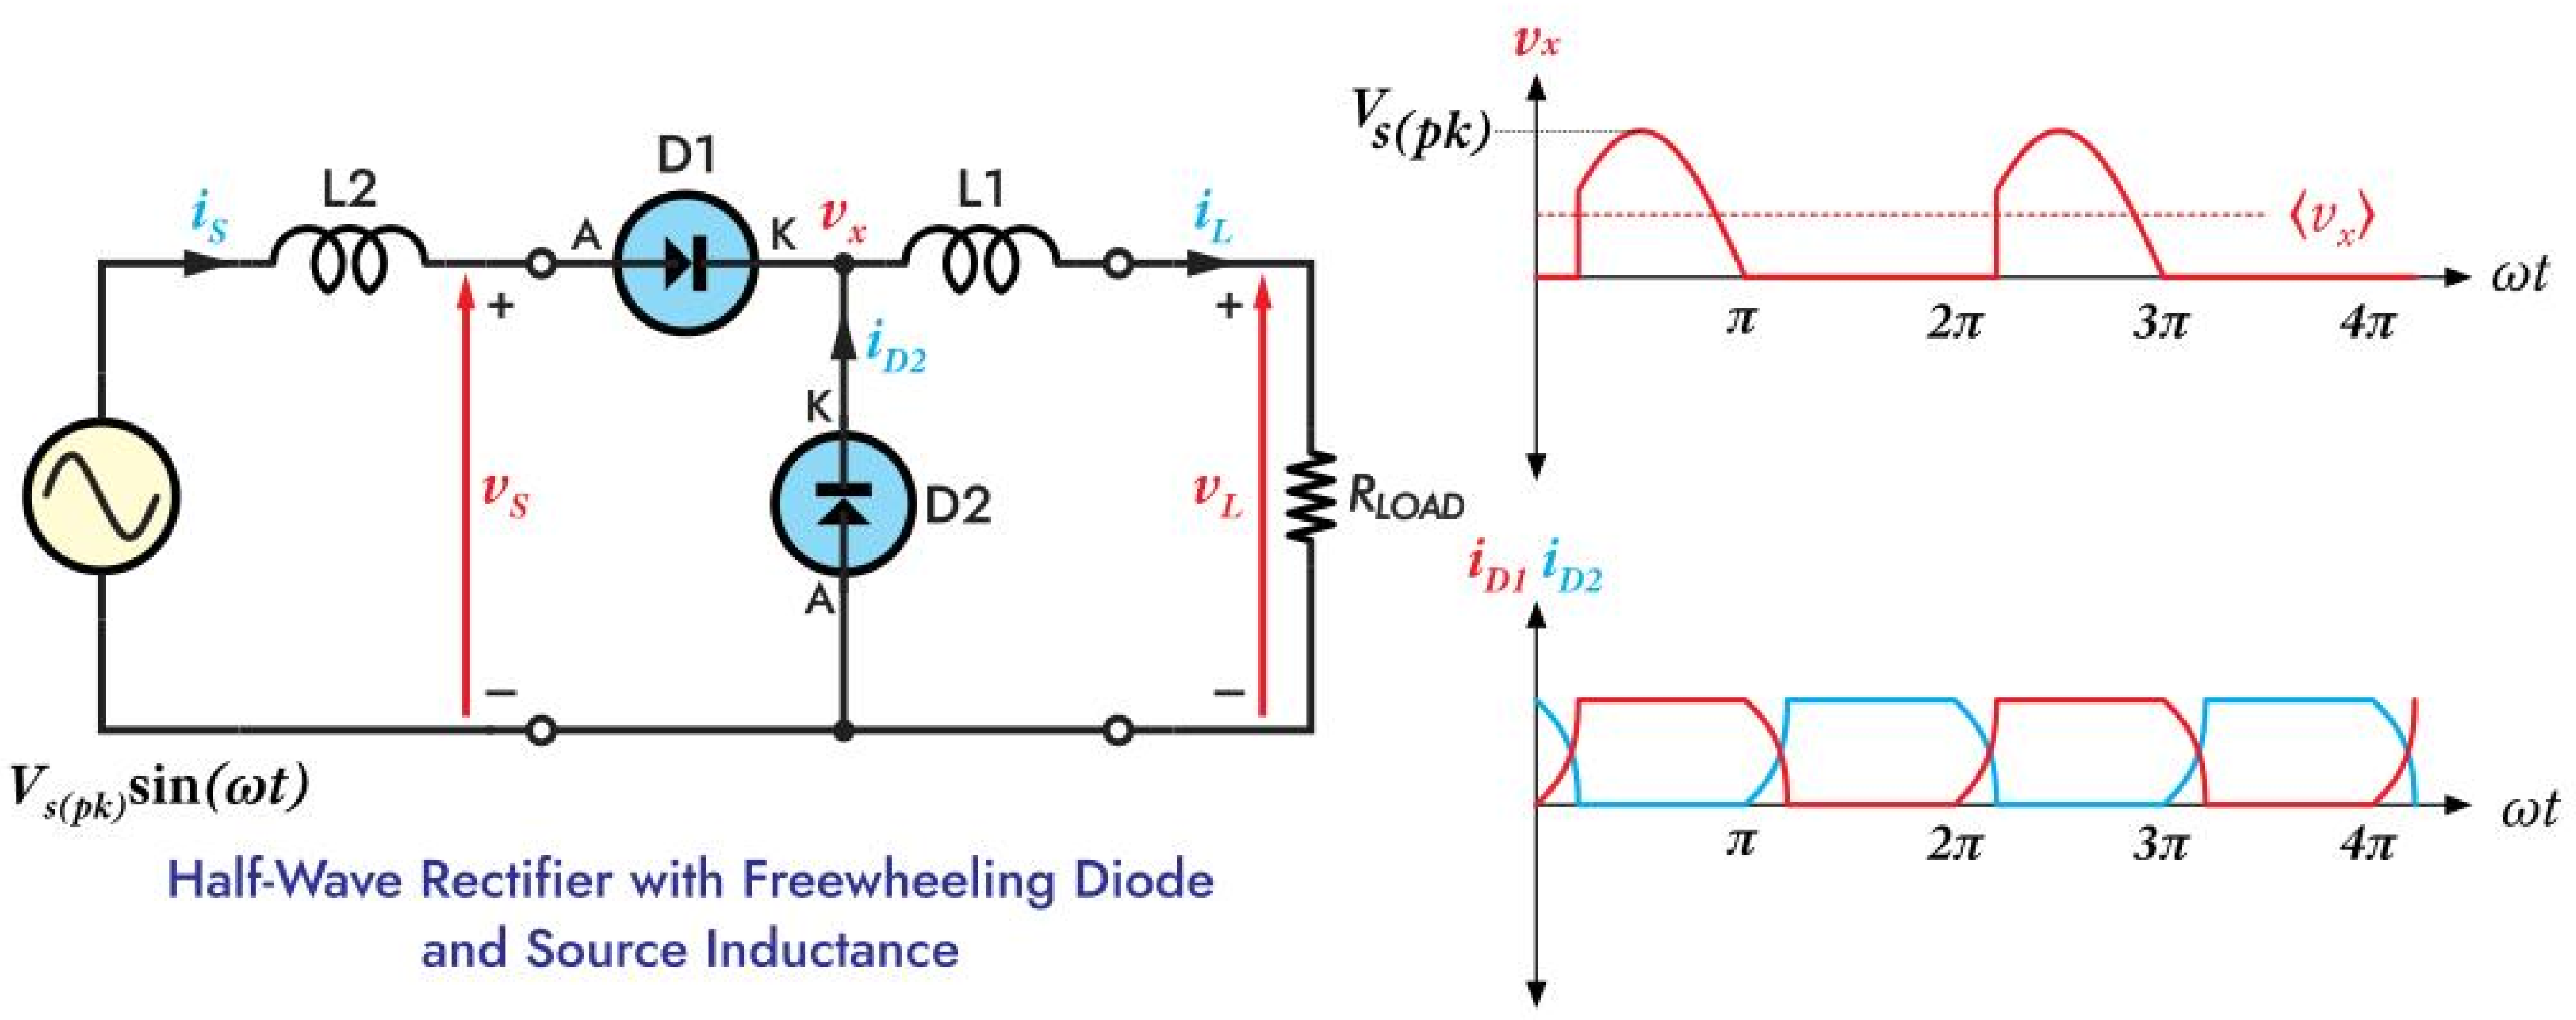
L2
D1
L1
i_S
A
K
v_x
i_L
i_{D2}
D2
v_S
v_L
RLOAD
$V_{s(pk)}\sin(\omega t)$
Half-Wave Rectifier with Freewheeling Diode and Source Inductance
v_x
$V_{s(pk)}$
$\langle v_x \rangle$
ωt
π
2π
3π
4π
i_{D1} i_{D2}

Fig.2: the non-zero source inductance L2 causes commutation, where both diodes conduct simultaneously for a short period, impacting power factor and regulation.

This circuit clearly will not work to create a smooth DC current in the load.

We have already seen the solution in the DC-DC examples; introduce a freewheeling diode, as shown in Fig.1(c). We can now make the inductor as large as we like, since the inductor current now has a path to flow when the source voltage reverses.

In a practical circuit, we would probably add a capacitor to the output, making an LC filter, but it can be ignored for the purpose of our analysis. After all, if the inductor is large enough to eliminate current ripple, no current will flow into or out of it.

So now the voltage at v_x will be a half-sinusoid, which we know from above has an average value of $V_{s(pk)} \div \pi$. Due to the periodic steady-state rules, the average load voltage will be the same. Ohm's law dictates that the average load current must be $V_{s(pk)} \div \pi R$, and since there is no current ripple, this is also equal to the DC current, I_L.

The input current will therefore have a rectangular shape, with an amplitude of I_L and a duty cycle of 50%, as shown in blue in the upper graph. The current though D2 (lower graph) will be similar, but out of phase by half a cycle.

Commutation

We have assumed a perfect voltage source up to this point, but we know that won't be the case in real life. If our rectifier is powered from the mains, it will see a source voltage with an inductive component.

If it comes from the mains via a transformer, there will be even more inductance, due to the transformer's leakage inductance. In fact, it is hard to think of a situation where there won't be any source inductance.

Adding source inductance to the circuit of Fig.1(b) makes no difference to its operation, because it is effectively in series with L1. However, things get more complicated if we add it to the single-phase rectifier with a freewheeling diode, as shown in Fig.2.

When D1 is conducting and D2 is off, the current through L2 is the DC load current, I_L, so the voltage across it is zero and v_x tracks the source.

However, a problem arises when the source voltage goes negative. D1 must remain conducting while current is flowing in L2, and the voltage v_x cannot go negative due to D2, so the voltage across L2 begins to rise until the energy stored in L2 is all transferred to the load via D1.

Therefore, the current through D1 does not drop instantly at the end of the half cycle; instead, it tapers off, as shown in the figure. This means D2 does not take over the current instantly either, and it ramps up in a complementary manner, because the total current flowing through L1 to the load remains fixed. D1 and D2 are therefore both in conduction for a (hopefully) short period.

The same thing happens in reverse when the positive half-cycle starts. The input current is now zero, so it takes a finite time for it to ramp up through L2 and D1 to the level of the load current. During this period, D2's current ramps down to maintain the constant load current.

While both diodes are conducting, the voltage at v_s is held at zero, truncating the beginning of the voltage half-cycle and slightly reducing the average output voltage.

This process, where current is transferred between the switches, is called commutation. The effect of commutation is to reduce the load regulation and change the shape of the source current.

The effect on regulation is equivalent to adding a resistor of value $X_{L2} \div 2\pi$ in series with the output. The effect on the input current is to make it more trapezoidal, meaning that the RMS source current for a given level of load current is higher with commutation than without.

Power factor

This segues nicely into the topic of power factor, which has become increasingly important in power electronics.

All the rectifiers we have looked at so far have a source current waveform that is non-sinusoidal. This means that the apparent power entering the rectifier, calculated as the source voltage times the source current, is higher than the real power delivered to the load, calculated as the average over a cycle of the instantaneous voltage times the instantaneous current.

How can this be? Consider the simple example in Fig.3. At the top, we have a sinusoidal source voltage (red trace) feeding some converter that produces a distorted (non-sinusoidal) current waveform (blue trace).

In this case, the source current i_S the sum of two sinusoidal currents: i_{S1}, at the same frequency as the source voltage, and i_{S2}, which is of lower magnitude but twice the frequency.

The lowest chart shows the instantaneous product of the source voltage v_s with the fundamental component of the current i_{S1} (in dark green) and the product of the source voltage v_s with the second harmonic component of the current i_{S2} (in light green).

The average power available from the fundamental component of the current is positive, but the average power available from the harmonic component is zero. This turns out to be true for all harmonics.

From Fourier theory, we know that any periodic current waveform can be decomposed into a series of sinusoids, including a fundamental component and its harmonics. However, if the source voltage is sinusoidal (like the mains), only the fundamental component of the current contributes any useful power to the load. This is the 'real power', which we designate ‹*p*›, in units of watts.

While the harmonic components of the current do not contribute to real power, they do contribute to the RMS value of the current, and therefore to the 'apparent' power – the product of RMS source voltage and RMS source current. We use S to describe apparent power, which has units of VA (volt-amps), and which will always be greater than or equal to the real power.

The ratio of real power to apparent power, ‹*p*› ÷ S, is the definition of power factor. It is a unitless quantity that varies between zero and one. A power factor of one means that the real and apparent power are equal, as would be the case for a resistive load, for example, and implies that the current is purely sinusoidal and in phase with the voltage.

A power factor less than one means that the real power doing useful work is lower than the apparent power being consumed.

Power factor is important because the electrical supply system has to be dimensioned for apparent power. For example, a UK domestic power outlet is rated to deliver 230V at 13A for a nominal apparent power of 2990VA. If you applied a unity power factor load (like a resistive heater), you can expect to get 2990W of usable power from such an outlet.

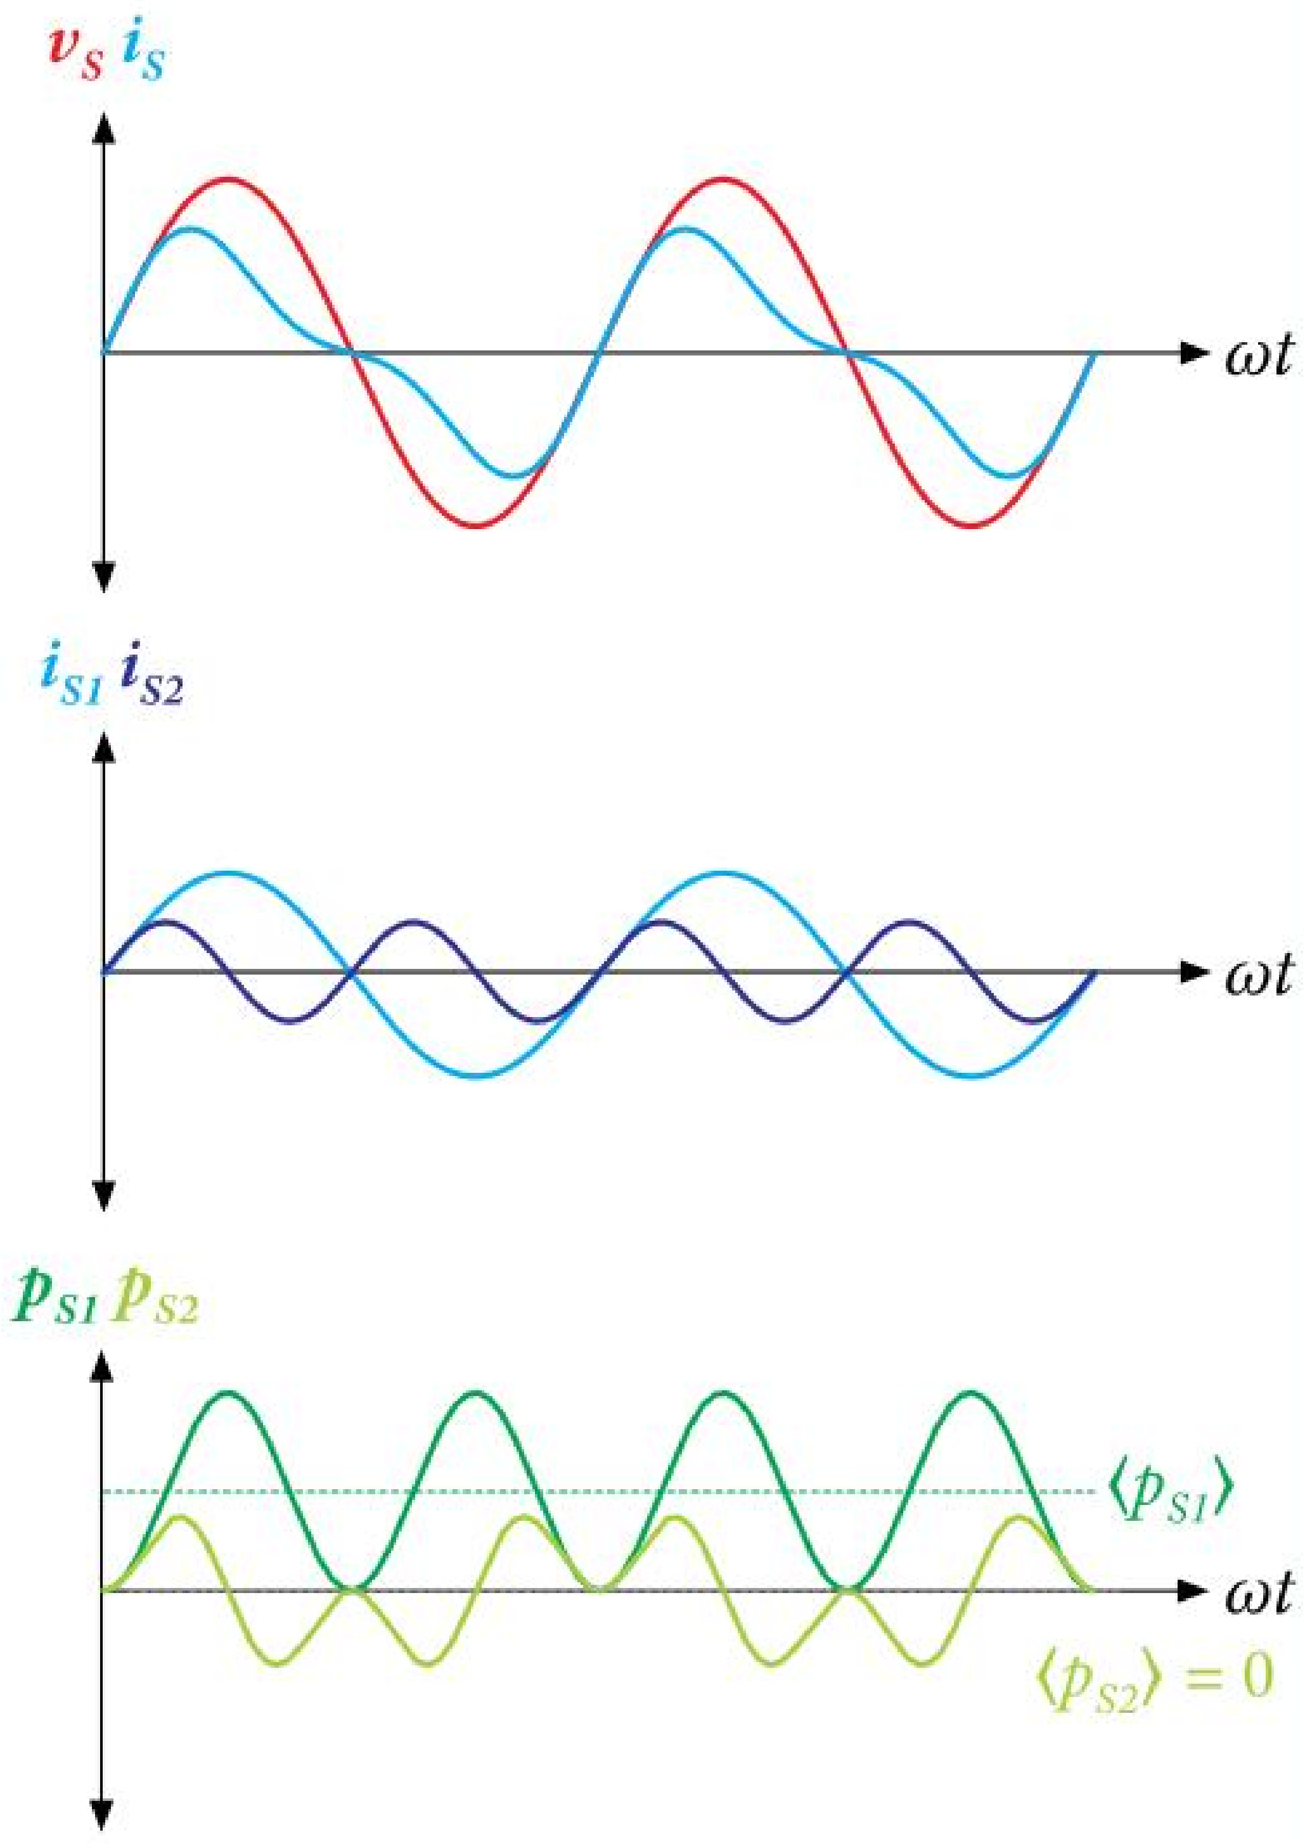

Fig.3: with a sinusoidal voltage source, only the fundamental component of the current waveform contributes to real power. The average power of any current harmonics is zero.

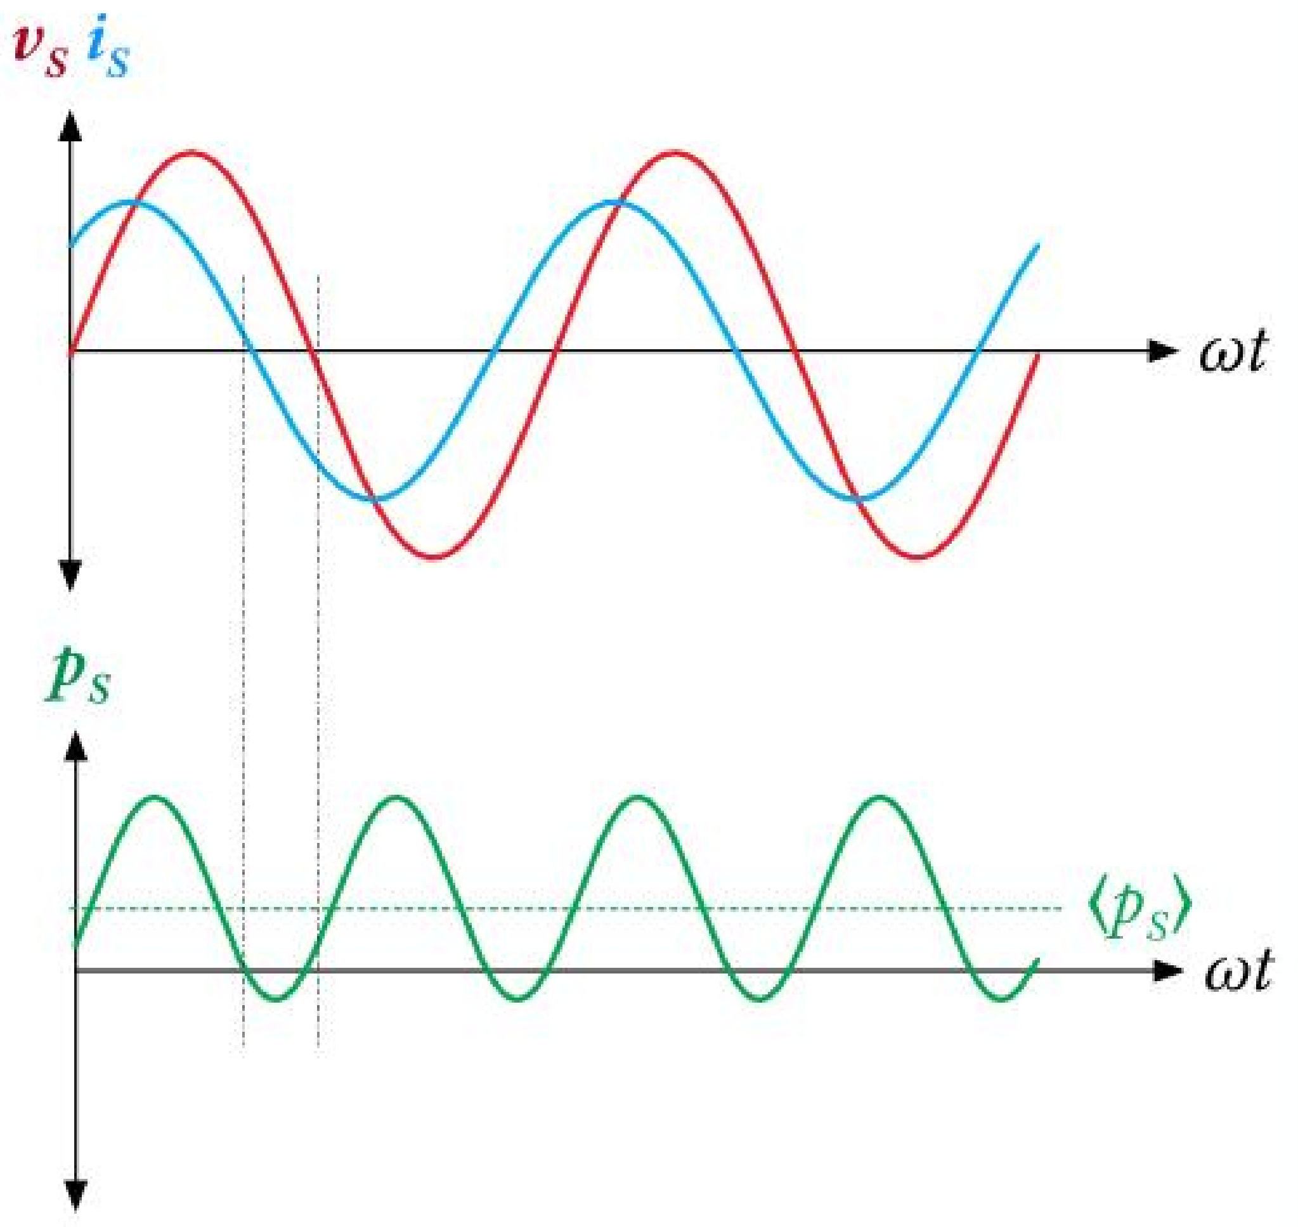

Fig.4: if the current is sinusoidal but out of phase with the voltage, the average power available will be reduced.

However, if the load produces a non-sinusoidal current, the usable power will be lower. If the power factor were 0.75, for example, you would only be able to get 2243W of useful power from the circuit, even though the RMS current sourced from the outlet would be 13A. Clearly, we can get the most out of the power distribution system by keeping the power factor high.

Mains-connected power electronics has become a major contributor to poor power factor in electricity distribution systems around the world. A variety of techniques are available to improve or 'correct' power factor.

We will cover some of these circuits in the next instalment of this series.

For completeness, I should point out that just eliminating current harmonics won't get you to unity power factor if the current is out of phase with the voltage. In fact, when I studied electrical engineering (many decades ago), the only discussion of power factor related to phase. The distortion component was skipped altogether because switch-mode supplies were far less common then.

Fig.4 shows why phase matters. A sinusoidal source voltage feeds a device that draws a current that is also purely sinusoidal, but slightly out of phase with the voltage. The green trace shows the instantaneous power obtained by multiplying the two.

If you compare this to the green power trace in Fig.3, you will see that instead of riding on the zero line, this trace is shifted down, reducing the average power compared to the in-phase case.

The vertical dashed lines show that as the phase shift increases, the zero-crossings of the voltage and current sinusoids diverge, so the power curve must drop to keep its zero crossings aligned. When either the voltage or current is zero, the instantaneous power must also be zero. If the phase shift reaches ±π/2 radians (±90°), the average power, and therefore the power factor, drops to zero.

If voltage and current are pure sinusoids, cos(ø), where ø is the phase shift, is a shorthand way to calculate power factor.

The power ratio equation is the one to use in power electronics as it works for both phase- and distortion-related power factor or any combination thereof.

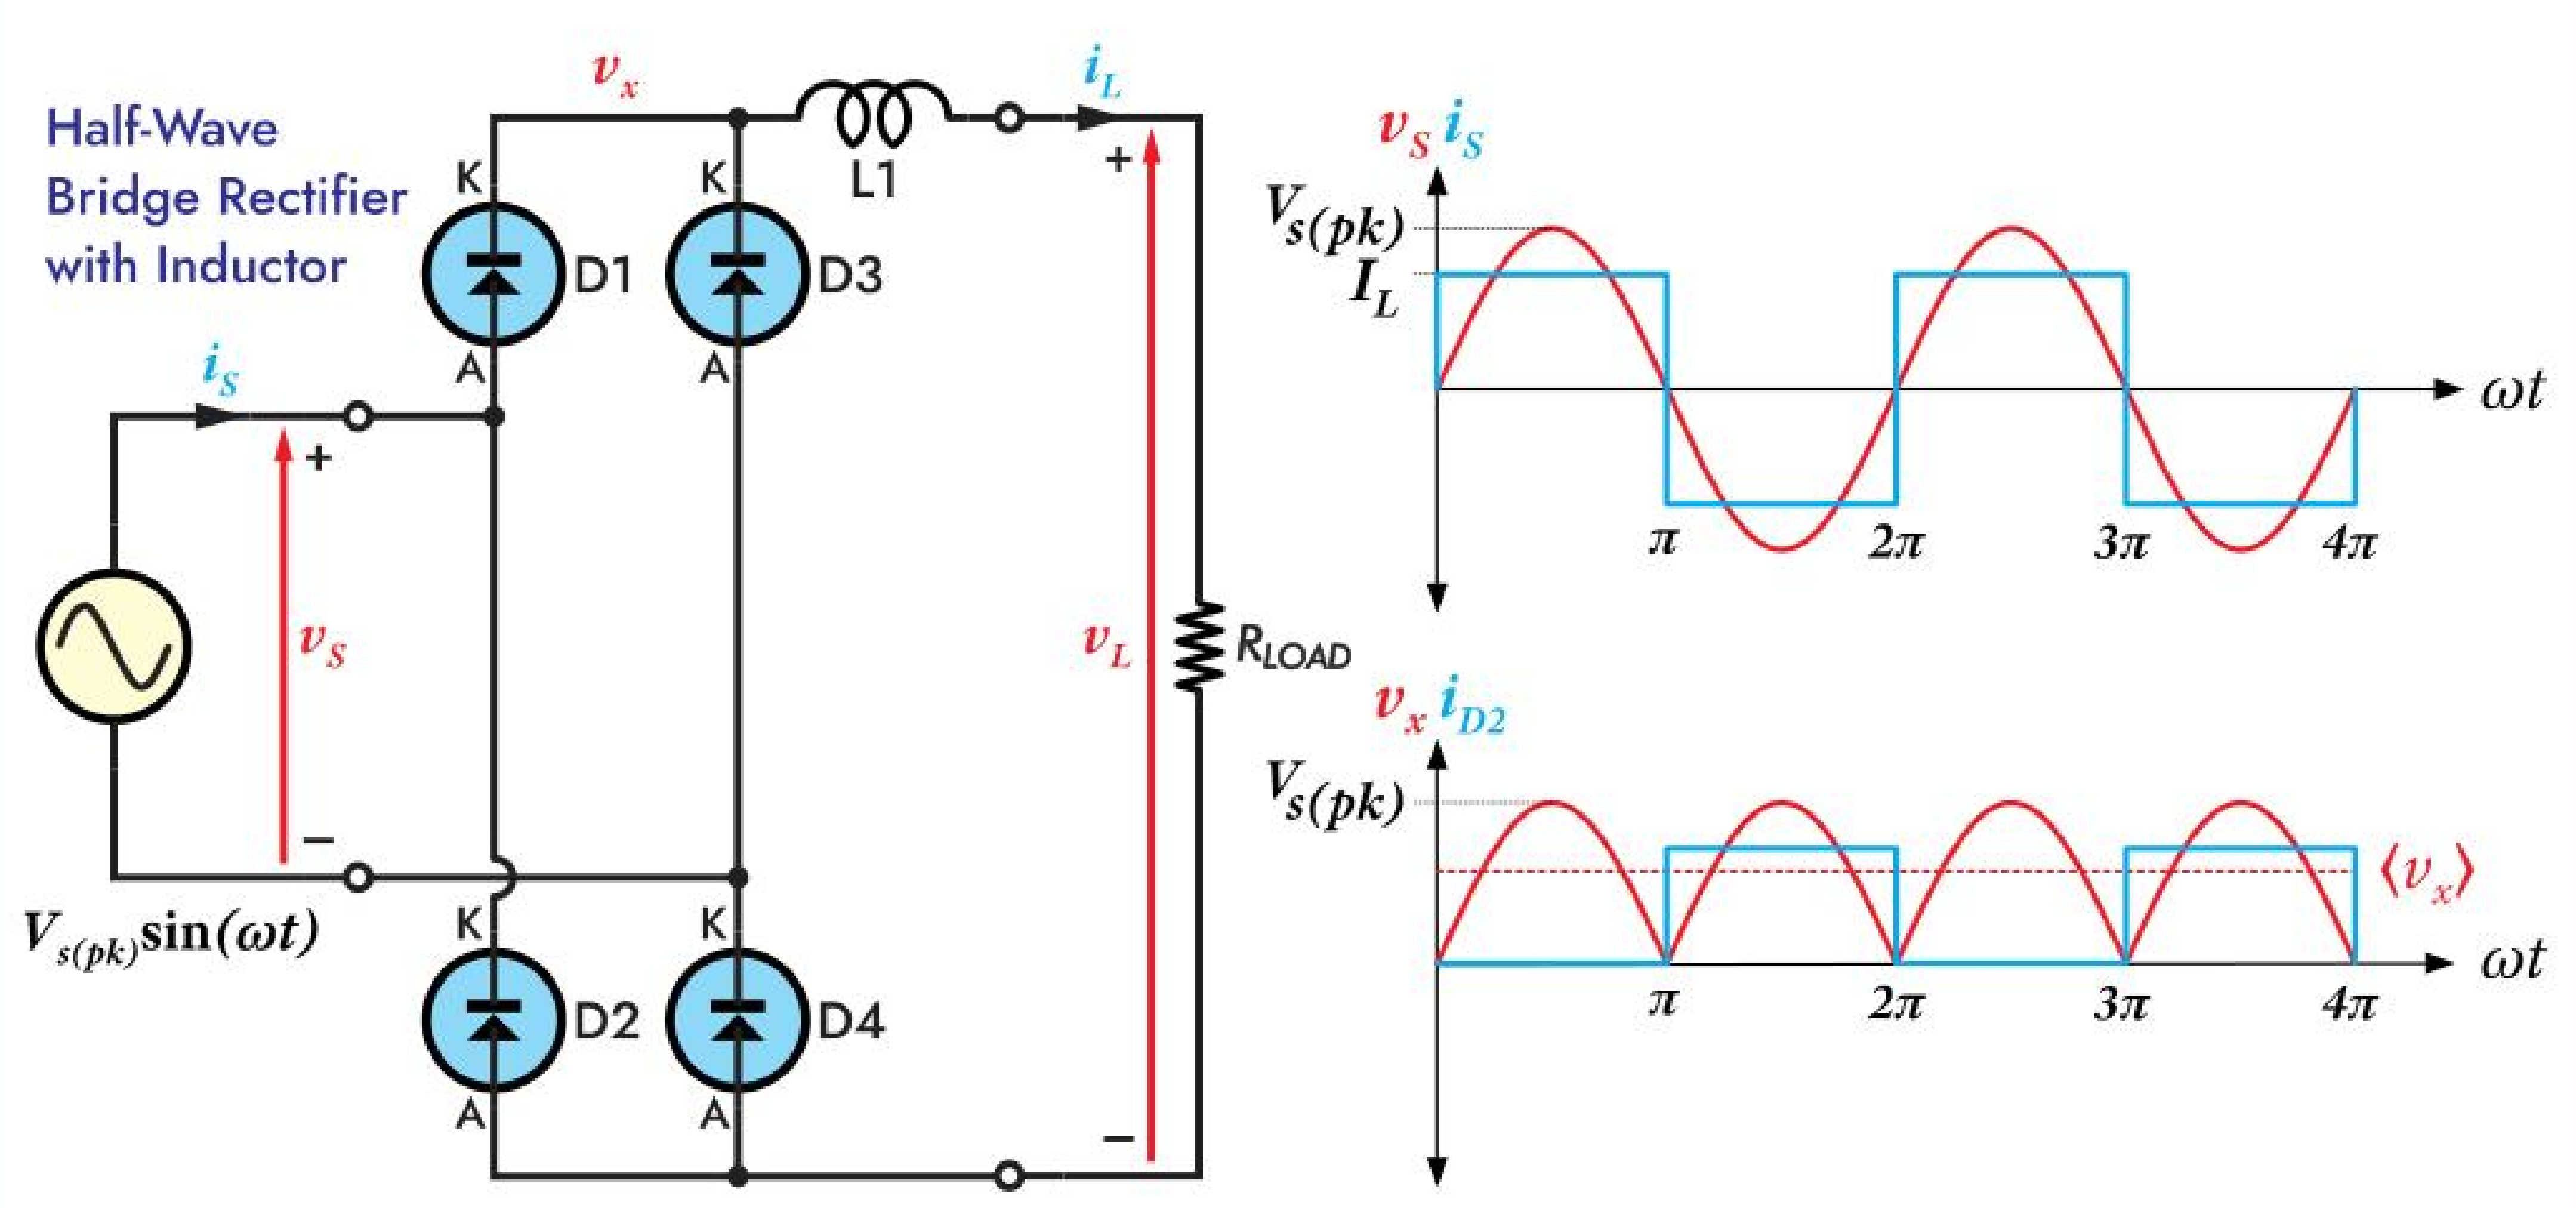

Fig.5: the full-wave bridge rectifier with inductor overcomes the disadvantages of half-wave rectifiers. The input current has no DC component, energy is delivered to the load on every half-cycle, and the power factor is much better.

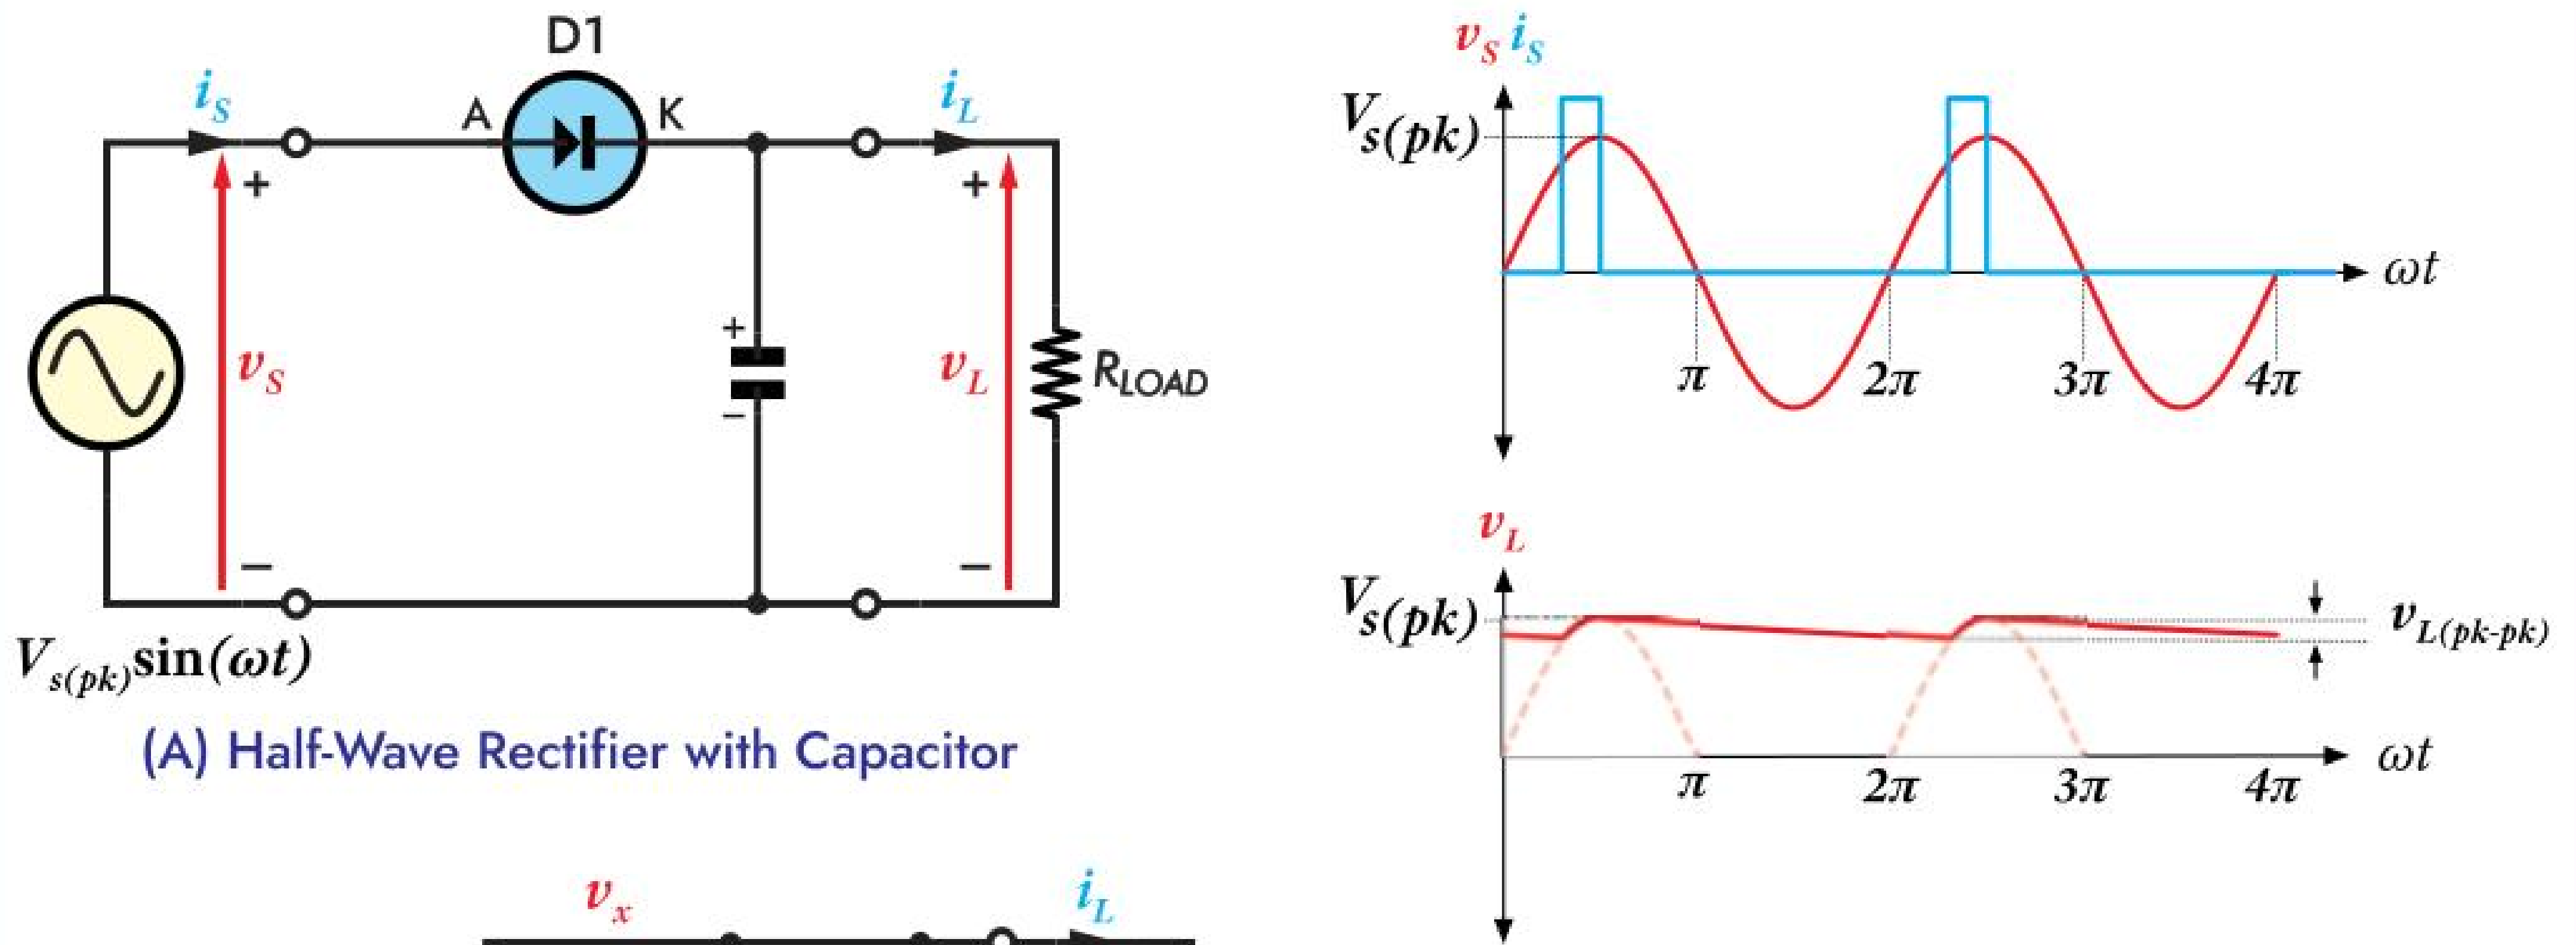

(A) Half-Wave Rectifier with Capacitor

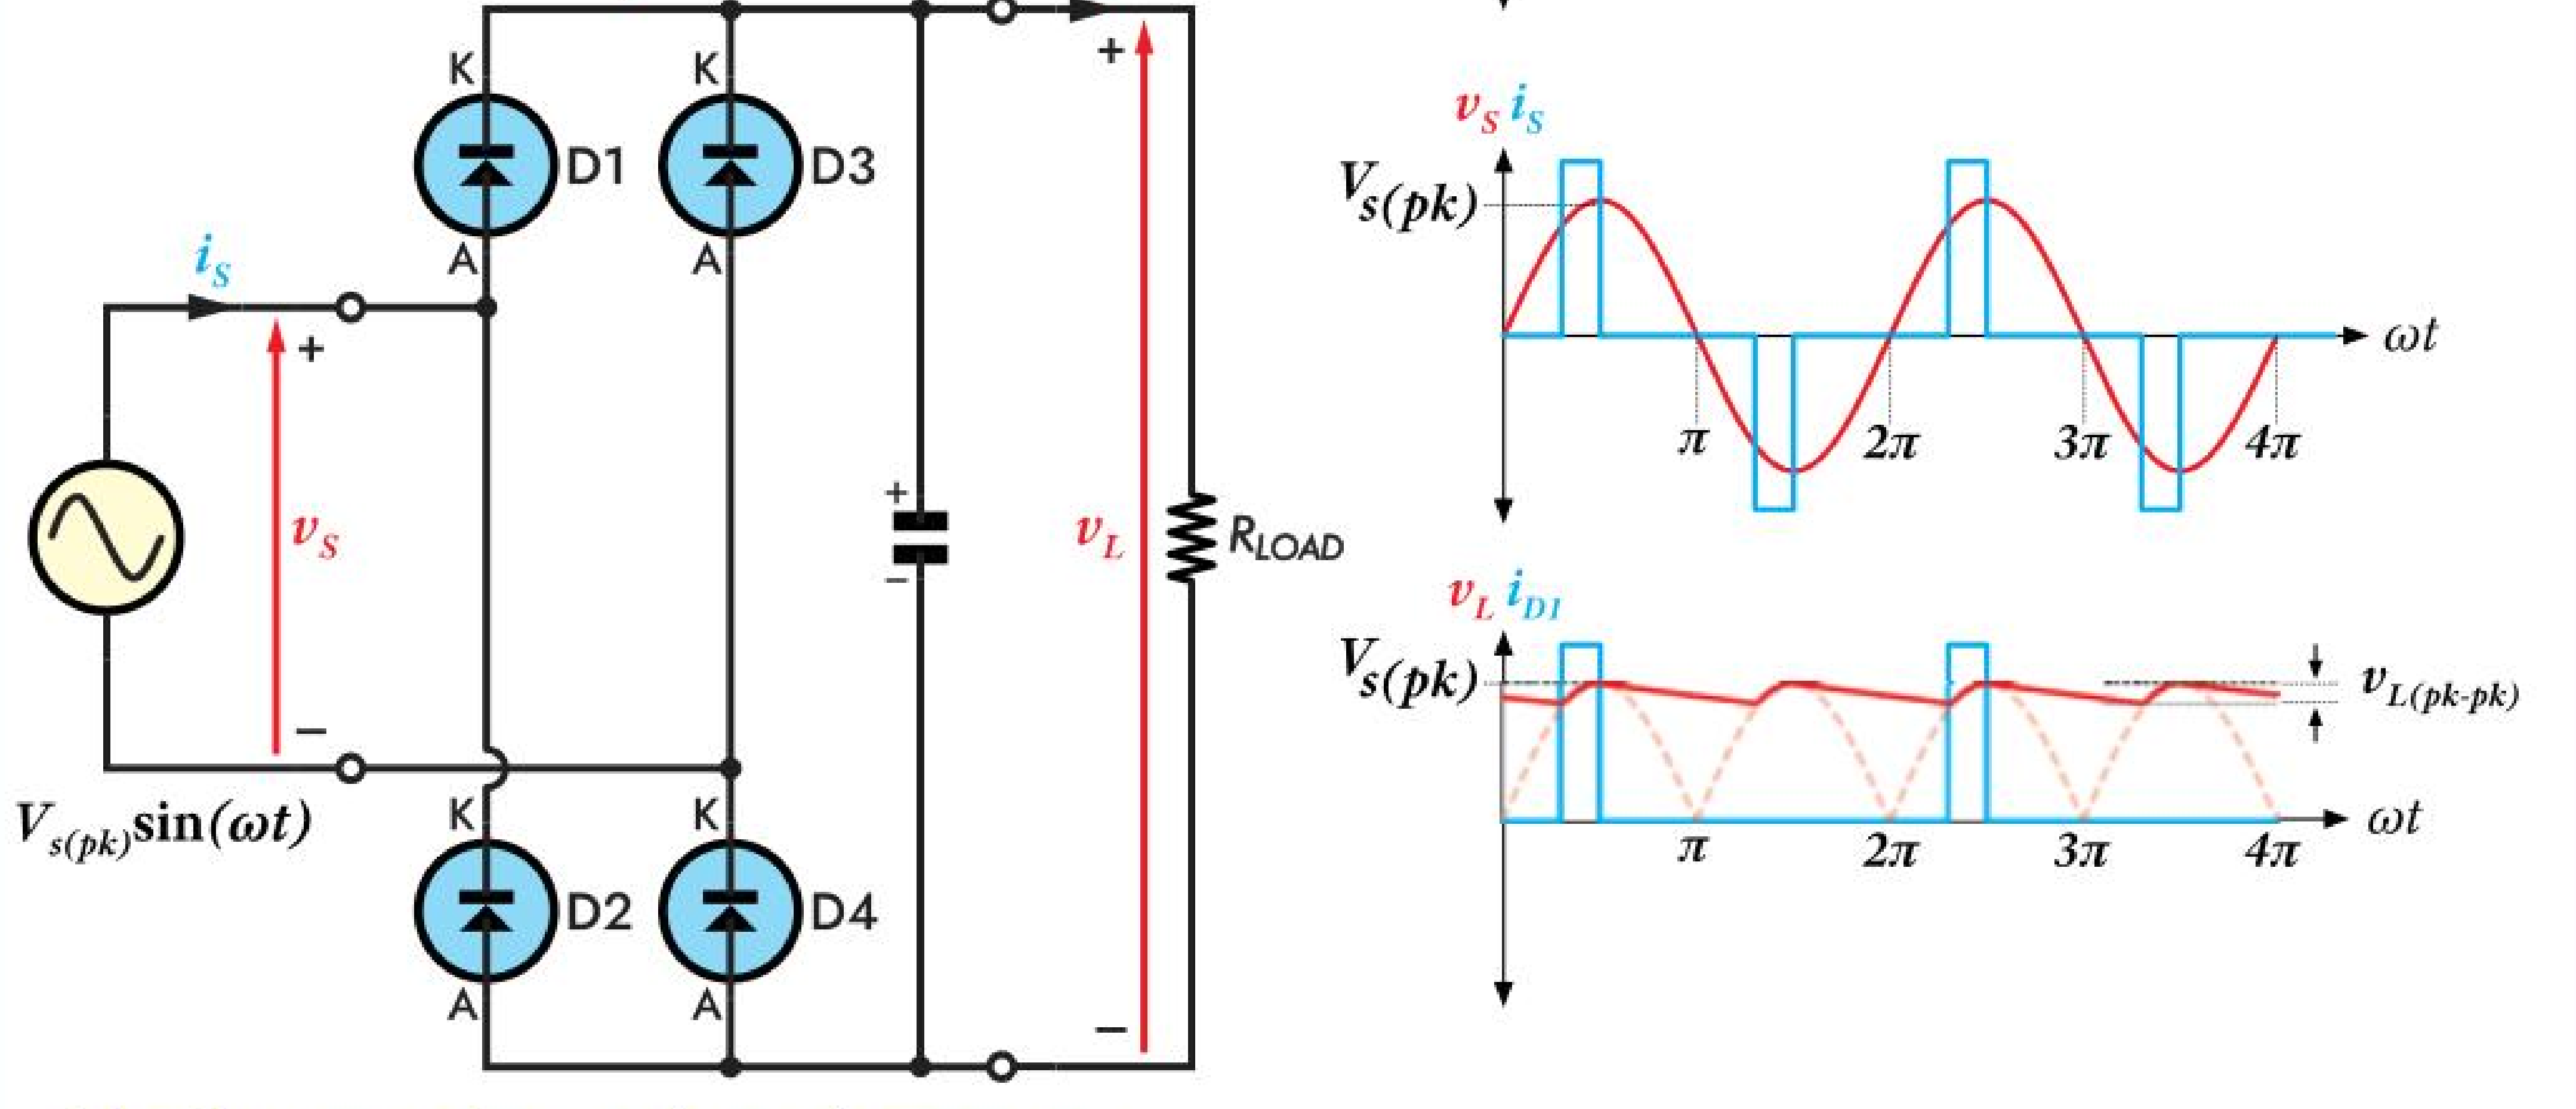

(B) Full-Wave Bridge Rectifier with Capacitor

Fig.6: we typically use a capacitor filter instead of an inductor. Circuit (A) is simple and cost-effective, but has a fairly poor power factor.

Full-wave rectifiers

The half-wave rectifiers that we described above are rarely used in practice since they suffer from three major drawbacks. First, the average input current is non-zero. This means there is a DC component to this current, which will not play nicely with transformers in the source network (and can accelerate conductor corrosion in some cases).

Secondly, they require large filtering components to achieve low voltage ripple because energy is supplied only during every other half-cycle. Lastly, they have a poor power factor.

We can confirm this pretty easily. Consider the half-wave rectifier with freewheeling diode in Fig.1(c). We saw that the average load voltage was $\langle v_L \rangle = V_{s(pk)} \div \pi$. We know the load current I_L is DC, so we can calculate the average power dissipated in the load (and therefore supplied by the source) to be $\langle p \rangle = V_{s(pk)} I_L \div \pi$.

The RMS input voltage is $V_{s(pk)} \div \sqrt{2}$ and the RMS input current is $\sqrt{I_L^2 \div 2} = I_L \div \sqrt{2}$, so the apparent power must be $S = V_{s(pk)} I_L \div 2$. Dividing $\langle p \rangle$ by S cancels the voltage and current terms, leaving a power factor for this topology of 2 ÷ π, which is about 0.64. Not great.

Fig.5 shows a full-wave bridge rectifier and its associated waveforms. You can think of this circuit as two single-phase rectifiers with freewheeling diodes – in positive half-cycles, D1 conducts and D3 is the freewheeling diode, while in negative half-cycles, D2 conducts and D4 is the freewheeling diode.

The full-wave rectifier addresses all the problems of the half-wave rectifier. The input current swings positive and negative alternately, so has an average value of zero and therefore no DC component. Power is supplied every half-cycle, doubling the output frequency and thus requiring less filtering to achieve a given level of voltage ripple. The power factor is also much better.

The average load voltage is twice that of the half-wave rectifier, so the average power is $\langle p \rangle = 2V_{s(pk)}I_L \div \pi$. The RMS input voltage is the same, but the RMS current is now just I_L, giving an apparent power of $S = V_{s(pk)} I_L \div \sqrt{2}$. Dividing $\langle p \rangle$ by S cancels the voltage and current terms, as before, but leaves us with a power factor of $2\sqrt{2} \div \pi$, which is about 0.90. This is much better.

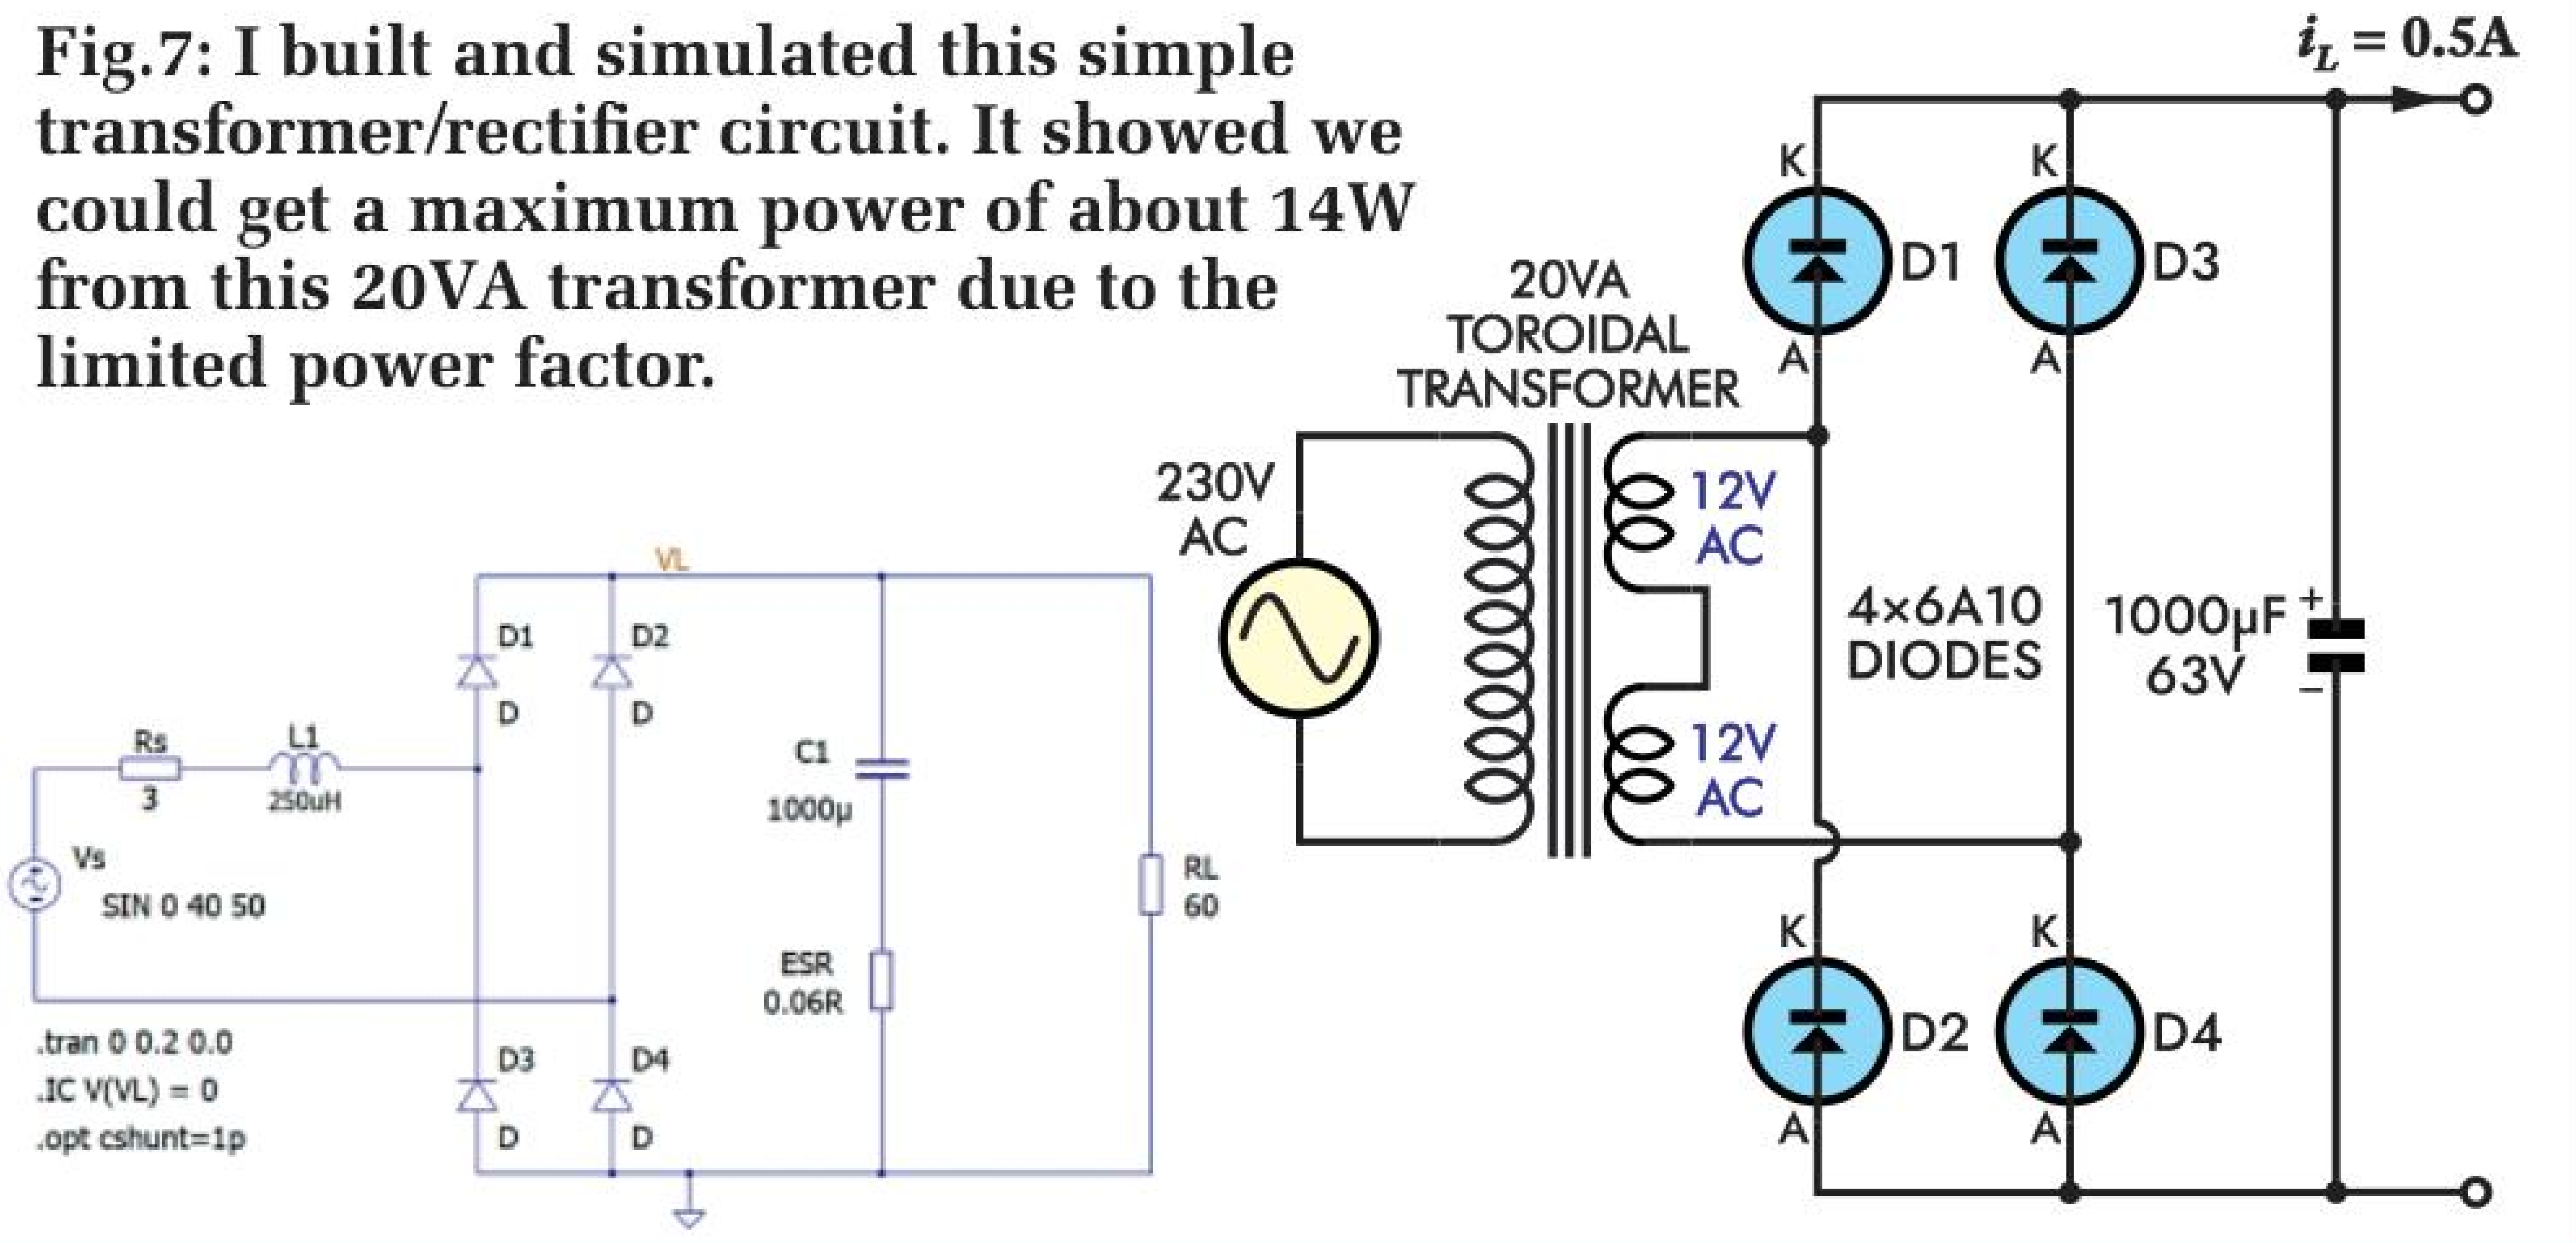
VL
D1
D2
D
D
Rs
3
L1
250uH
C1
1000μ
Vs
SIN 0 40 50
RL
60
ESR
0.06R
.tran 0 0.2 0.0
.IC V(VL) = 0
.opt cshunt=1p
D3
D4
D
D
i_L = 0.5A
K
K
D1
D3
A
A
20VA
TOROIDAL
TRANSFORMER
230V
AC
12V
AC
12V
AC
4×6A10
DIODES
1000μF
63V
+
K
K
D2
D4
A
A

Fig.7: I built and simulated this simple transformer/rectifier circuit. It showed we could get a maximum power of about 14W from this 20VA transformer due to the limited power factor.

Capacitive filters

Of course, all of this assumes the presence of a large inductor to smooth the current, but this is not how we usually build rectifier-filter circuits. For the most part, we simply add a capacitor directly after the half- or full-bridge, as shown in Fig.6.

The circuit acts like a peak detector, with the capacitor charging to $V_{s(pk)}$ via the diodes at the crest of the half- or full-wave rectified waveform (shown dotted). The capacitor discharges via the load until the next peak.

The voltage ripple can be (roughly) approximated by assuming that the voltage takes on a sawtooth profile (ie, it charges instantly at the crest and that the discharge is linear). For a half-wave rectifier, this gives $V_{L(pk\text{-}pk)} = I_L \div fC$, where f is the source frequency and C is the capacitance. For a full-wave rectifier, the ripple is half of this, ie, $V_{L(pk\text{-}pk)} = I_L \div 2fC$. This approximation is quite pessimistic for small supplies where the source impedance is relatively high, as we shall see.

This topology means that current only flows into the capacitor for a short period, resulting in a current waveform that is made up of narrow spikes of current aligned with the crests of the input voltage. The width of the current spikes is related to the ripple (the lower the ripple, the narrower the spikes), the source impedance and the capacitor's ESR.

All of this is difficult to calculate, but is a great candidate for simulation and experimentation.

A practical example

I had a 20VA, 240V to 12+12V toroidal transformer in my junk box, so I decided to build the simple full-bridge AC-to-DC converter shown on the right-hand side of Fig.7 to see how it performed. Notice that the transformer is specified for apparent power.

The transformer windings are connected in series to get a nominal 24V RMS, which is rectified by four chunky 6A10 (6A, 1000V) diodes I had lying around, and filtered by a 1000μF 63V electrolytic capacitor.

First, I measured the open-circuit voltage of the transformer (28.6V RMS), the DC resistance of the secondary windings (1.3Ω each), the transformer leakage inductance (250μH) and the filter capacitor's ESR (0.06Ω). This allowed me to build the simulation model shown in Fig.7. The simulation results are shown in Fig.8.

The average output voltage is 31.6V, with a peak-to-peak ripple of 3.4V. The average output power is therefore 16.0W. The input current is shaped as we would expect, but the input voltage shows a flattened top. This is due to the voltage drop across the source impedance when the current pulses occur, and is typical for supplies of this size.

The simulator calculated the RMS input voltage and current to be 26.2V and 0.95A, respectively, for an apparent power of 24.9VA (a little higher than our transformer's rating). The power factor is therefore 0.64.

The simulation compares well with the measured results below. The average load voltage is 30.6V and there is 4.8V peak-to-peak ripple. The output power is therefore 15.3W.

These measurements were taken

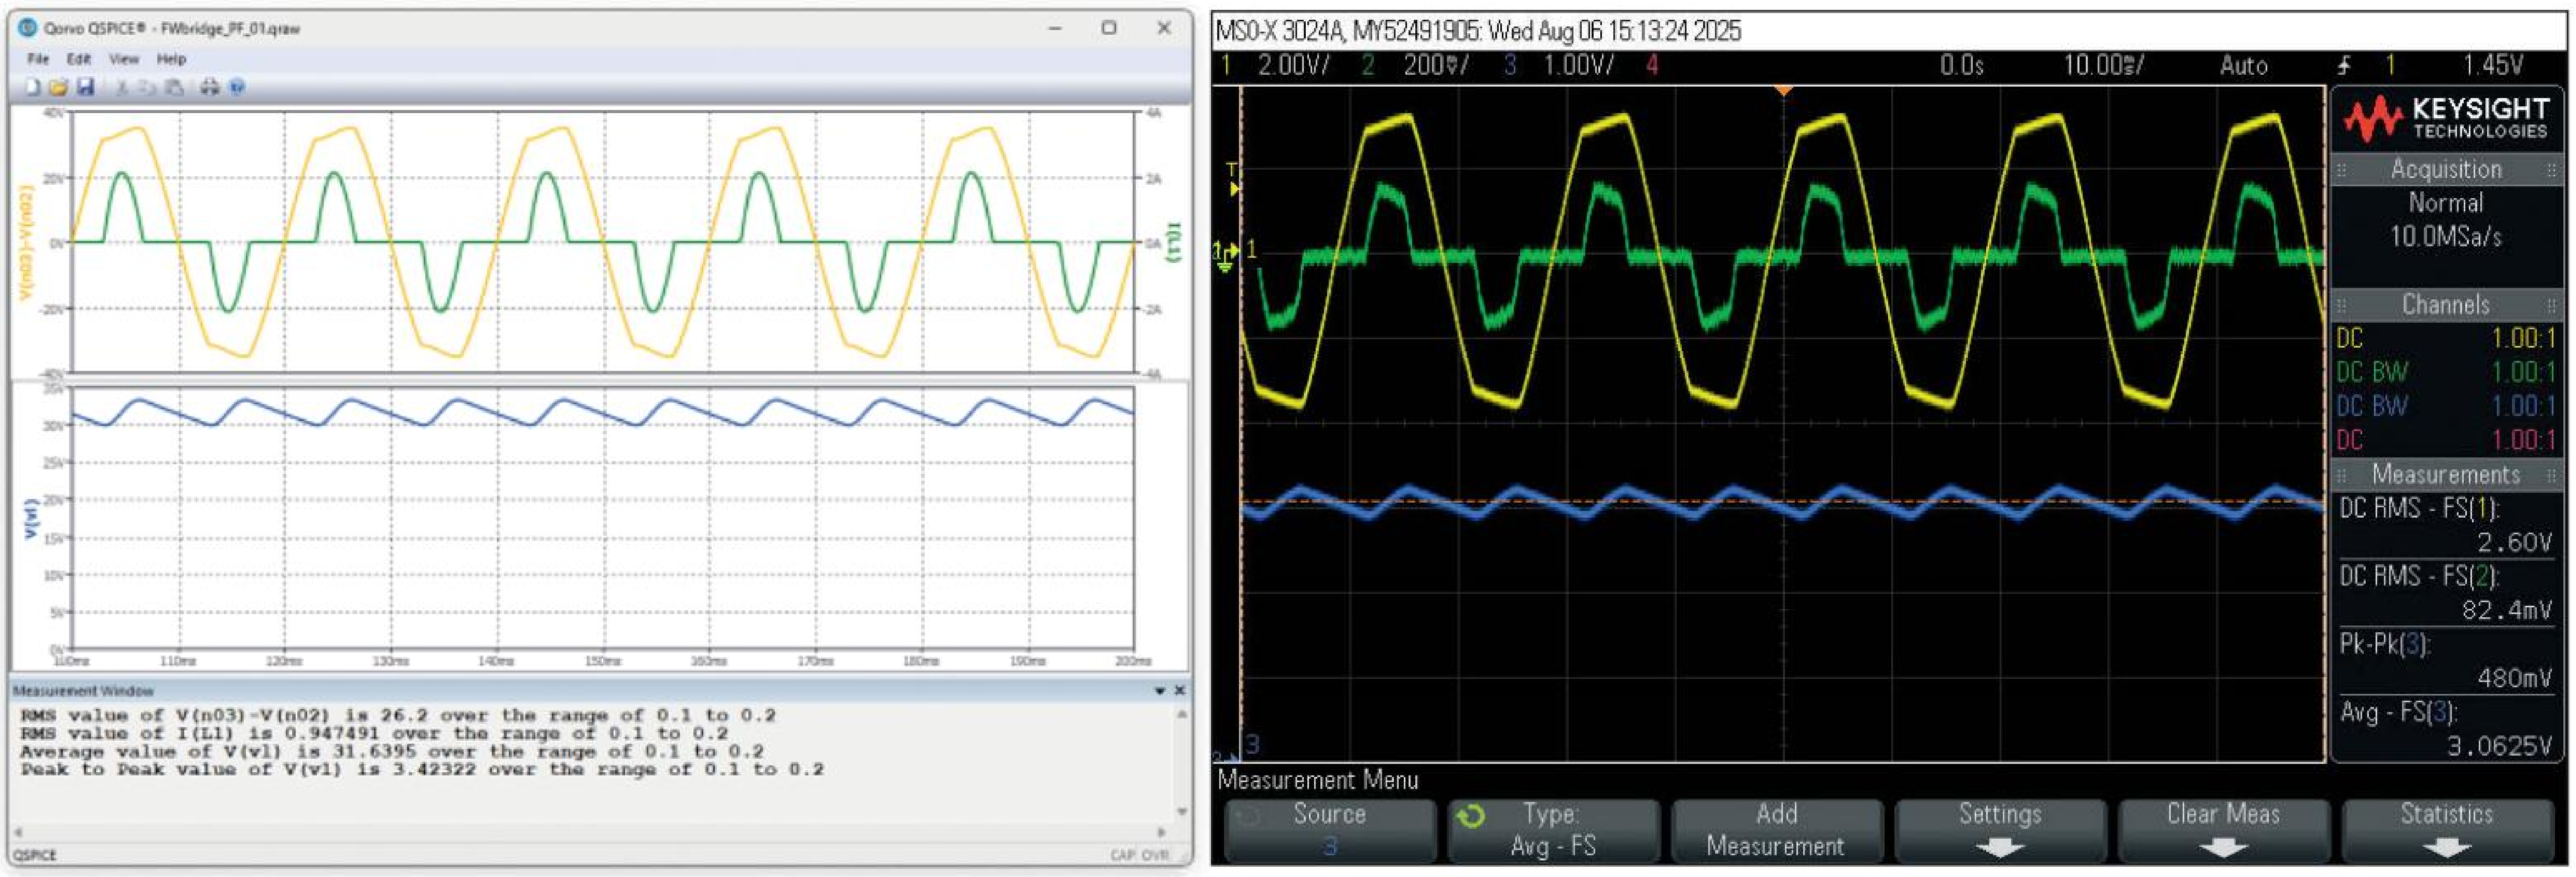
Qorvo QSPICE® - FWbridge_PF_01.qraw
File Edit View Help
V(n03)-V(n02)
I(L1)
V(vl)
Measurement Window
QSPICE
CAP OVR
MS0-X 3024A, MY52491905: Wed Aug 06 15:13:24 2025
1 2.00V/ 2 200mV/ 3 1.00V/ 4
0.0s 10.00ms/ Auto 1 1.45V
KEYSIGHT TECHNOLOGIES
Acquisition
Normal
10.0MSa/s
Channels
DC 1.00:1
DC BW 1.00:1
DC BW 1.00:1
DC 1.00:1
Measurements
Measurement Menu
Source
3
Type:
Avg - FS
Add Measurement
Settings
Clear Meas
Statistics

```
RMS value of V(n03)-V(n02) is 26.2 over the range of 0.1 to 0.2
RMS value of I(L1) is 0.947491 over the range of 0.1 to 0.2
Average value of V(vl) is 31.6395 over the range of 0.1 to 0.2
Peak to Peak value of V(vl) is 3.42322 over the range of 0.1 to 0.2
```

DC RMS - FS(1): 2.60V
DC RMS - FS(2): 82.4mV
Pk-Pk(3): 480mV
Avg - FS(3): 3.0625V

Fig.8: the experimental results agree fairly well with those obtained by simulation.

using a Current Probe and two Differential Probes, described in the November 2025 and January 2026 issues of this magazine, respectively, so the appropriate scaling factors need to be applied.

The RMS input voltage and current are 26.0V and 0.824A for an apparent power of 21.9VA (still a touch too high for the transformer in the long term). The power factor is therefore 0.69, slightly better than the simulation, but nothing to get excited about.

The important thing to note here is that the relatively low power factor puts an upper limit on the real power you can get from a given transformer. For this 20VA toroid, it is about 14W.

I should also point out that this type of rectifier/filter results in fairly high 100Hz current ripple in the capacitor, which raises its internal temperature and potentially shortens its life. In this circuit, the capacitor ripple current is 0.8A RMS. Electrolytic caps usually come with a maximum 100Hz ripple current specification, so it is worth checking you are not exceeding this limit.

The ripple current rating is one of the reasons you almost always see large filters made up of multiple parallel capacitors. If the capacitors are identical the ripple current rating of the bank is the sum of the ripple current rating of each capacitor.

5-year collections
2019-2023
£49.95
2018-2022
£49.95
2017-2021
£49.95
2016-2020
£44.95
Purchase and download at:
www.electronpublishing.com

The ‘peaky’ current waveform has an impact on the output voltage you will achieve with this circuit. There are two diode drops between the peak value of the transformer secondary voltage and the voltage across the filter capacitor.

These will likely be higher than the nominal 0.6-0.7V you might expect because the capacitor only charges when the current is at its peak.

The diode data sheet should provide a curve called “instantaneous forward characteristic” or similar, which relates forward voltage drop to peak forward current. In the example of the 6A10 diodes I used, this curve shows a forward voltage of 0.8V at the peak current we see in the simulation, for a total drop of 1.6V.

It is not at all unusual for the voltage drop to approach 2V if larger currents are involved. This drop can eat up a significant portion of the available voltage in low-voltage applications. As an aside, this is why active rectifiers are becoming more popular (see our November 2024 & September 2025 designs); they involve very little voltage loss and so improve efficiency.

Inrush current

This topology also comes with potential inrush current concerns. When power is first applied, and the capacitors are fully discharged, the inrush current is limited only by the supply impedance and the capacitor ESR.

This is rarely a problem with low-voltage supplies fed by relatively small transformers like this one, but can be a big problem for off-line rectifier/filters and very large capacitor banks, as you might find in a high-power audio amplifier, for example.

The Variable Speed Drive for Induction Motors (December 2025 & January 2026) used a bank of five 330μF 400V capacitors to filter the rectified mains.

A simulation made at the time showed that without inrush limiting circuitry, the peak inrush current would be around 200A, almost certainly tripping the supply circuit breaker and maybe damaging the rectifier diodes.

In that case, we used a special inrush-limiting thermistor with a cold resistance of 10Ω to limit these peaks to less than 35A. The thermistor’s resistance drops as current passes through it, and the model we used is rated to conduct 15A continuously.

These low-cost inrush limiters are available in various sizes and packages and are very common in off-line converters of all sizes. There was even one in the flyback converter we looked at last month.

You can use a resistor to limit inrush if you short it out with a relay or similar after it has done its job, but you need to be sure that the resistor can withstand the short pulse of power that occurs during inrush. In the case of the Variable Speed Drive, this peak power is well over 100W for a few milliseconds. Some power resistors are specified for pulse power, but many are not, so be careful.

Other topologies

There are many variants of the full-wave rectifier, and a couple of the more common ones are shown in Fig.9. The load voltages shown assume the diode voltage drops are negligible.

At the top is a centre-tapped variant that is a little more efficient for low-voltage supplies than the full bridge, since the current passes through only one diode instead of two. This comes at the expense of a more complex transformer, but in reality, dual secondary windings are common in small off-the-shelf transformers.

The middle rectifier uses a dual-winding transformer to produce a symmetric split (±) power supply – very useful for audio amplifiers or op amp circuits.

The final circuit is a full-wave voltage doubler that is effectively two half-wave rectifiers in series. During positive half-cycles, the upper capacitor is charged almost to the peak of the secondary voltage, and during negative half-cycles, the lower capacitor is charged to a similar voltage. The result is an output voltage twice what could be expected from the same transformer with a full bridge rectifier.

Phase control

No discussion of rectifiers would be complete without introducing phase-controlled rectifiers. This is a technique that is used less these days than it used to be, but is still relevant in industrial applications where very high currents must be controlled.

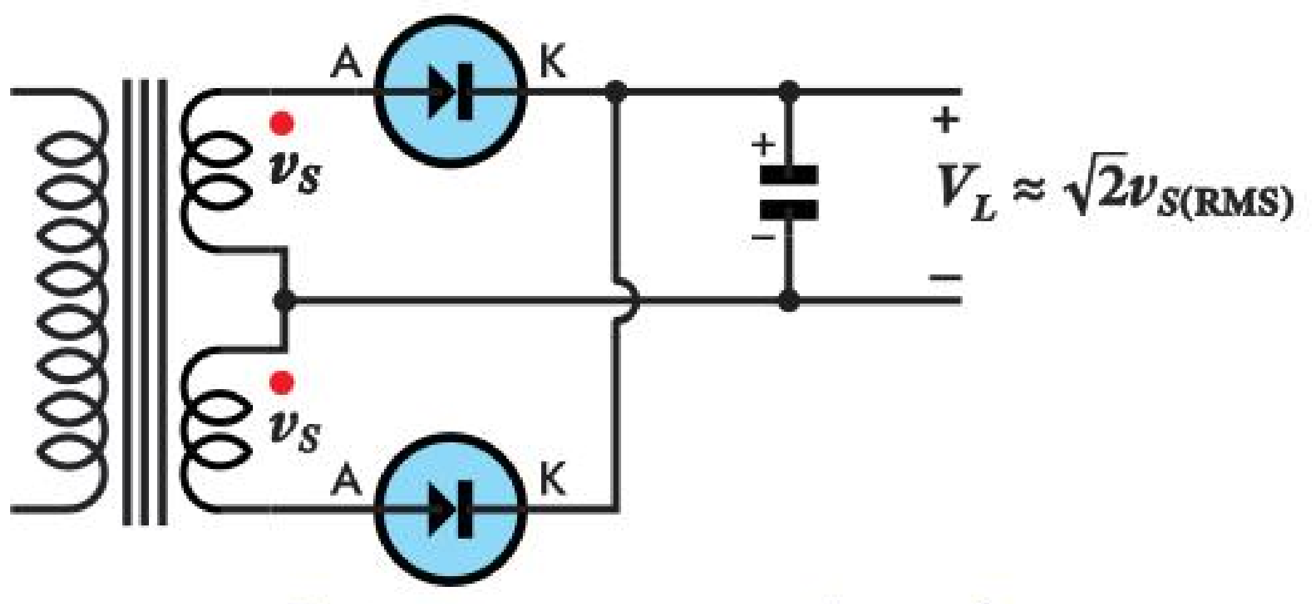
A
K
v_S
v_S
A
K
+
−
$V_L \approx \sqrt{2}v_{S(RMS)}$
+
−

Full-Wave Centre-Tapped Rectifier

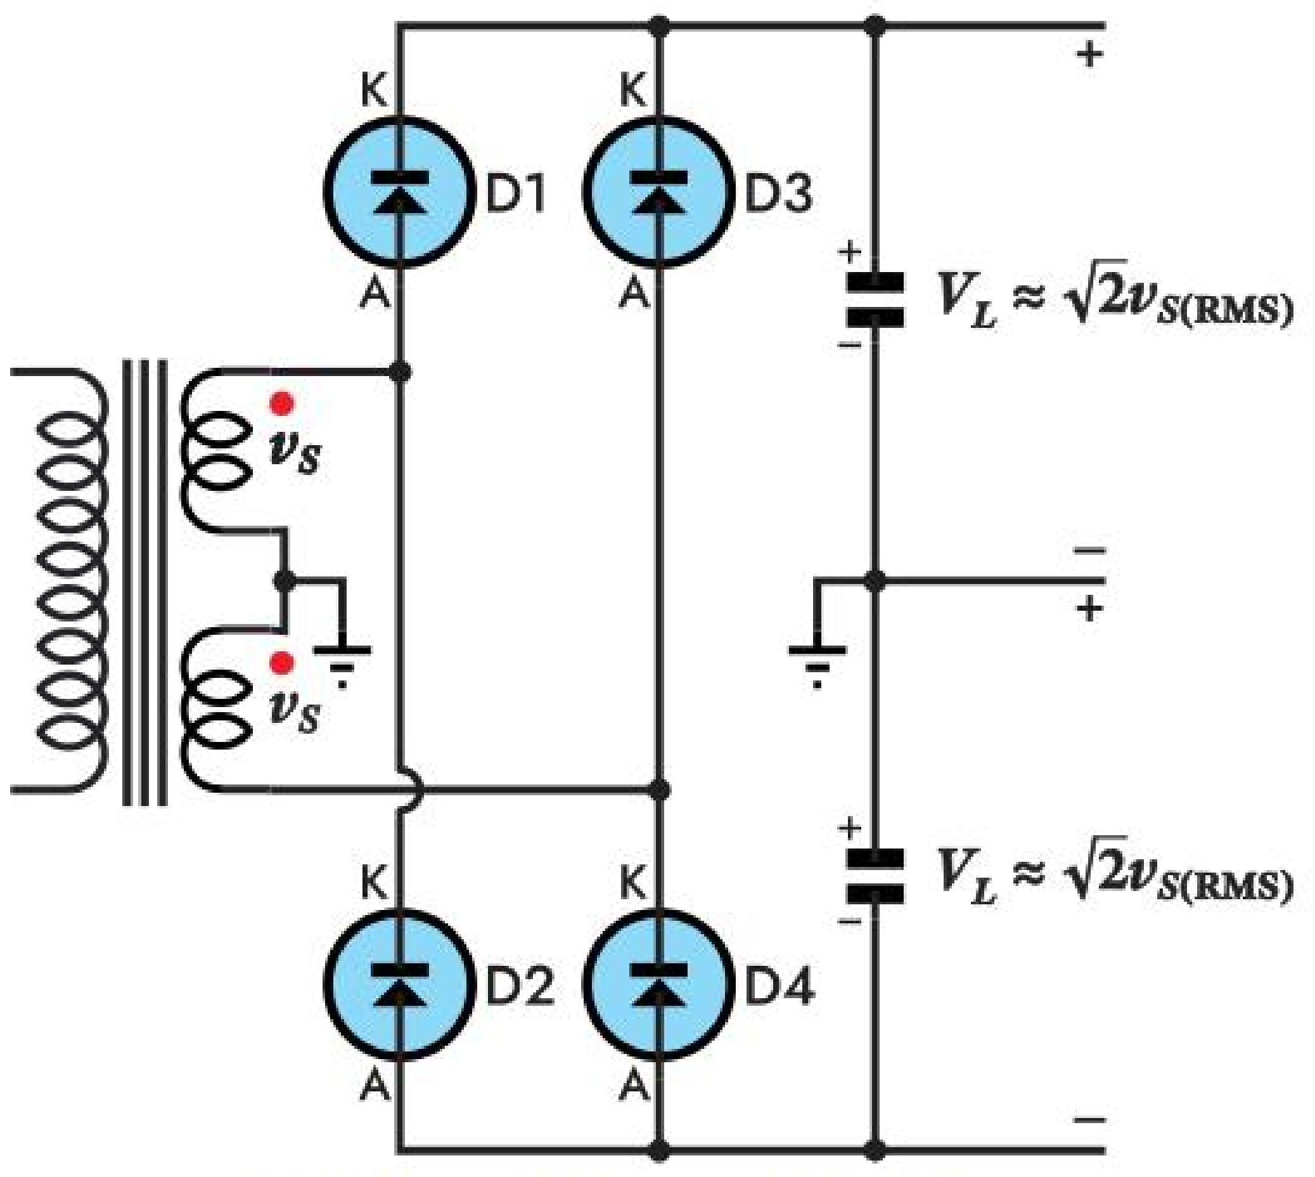
K
D1
A
K
D3
A
v_S
v_S
K
D2
A
K
D4
A
+
−
$V_L \approx \sqrt{2}v_{S(RMS)}$
+
−
$V_L \approx \sqrt{2}v_{S(RMS)}$
+
−
+
−

Full-Wave Split Supply Rectifier

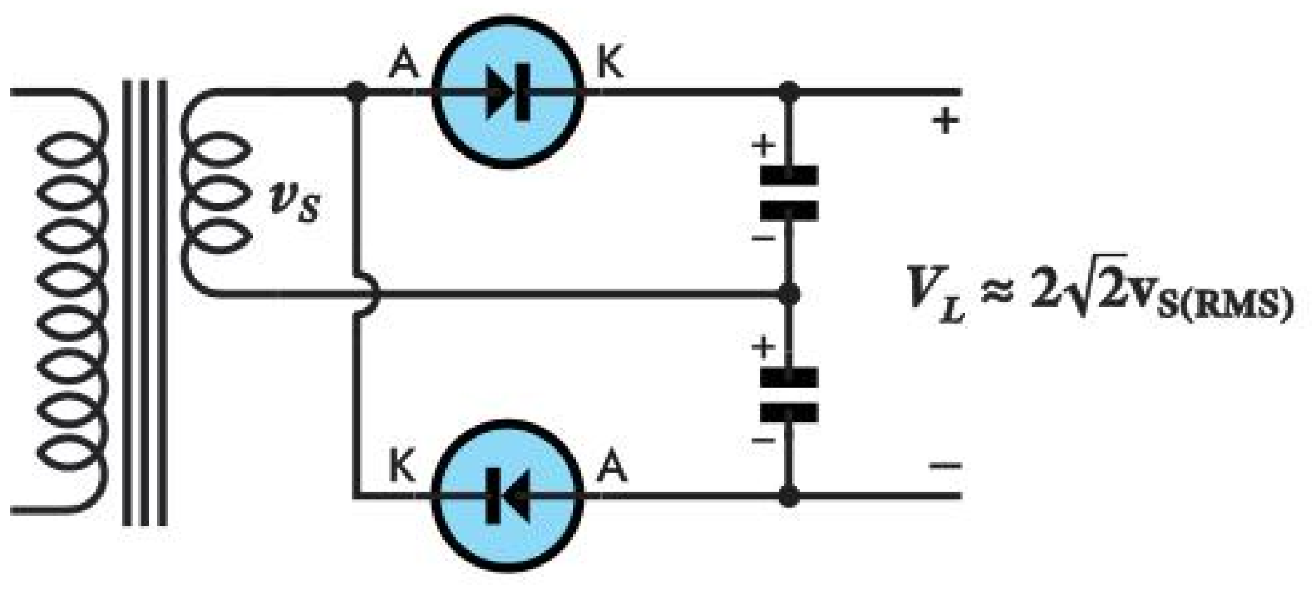
A
K
v_S
K
A
+
−
+
−
$V_L \approx 2\sqrt{2}v_{S(RMS)}$
+
−

Full-Wave Voltage Doubler

Fig.9: here are some variations on the full-wave rectifier theme that might come in handy.

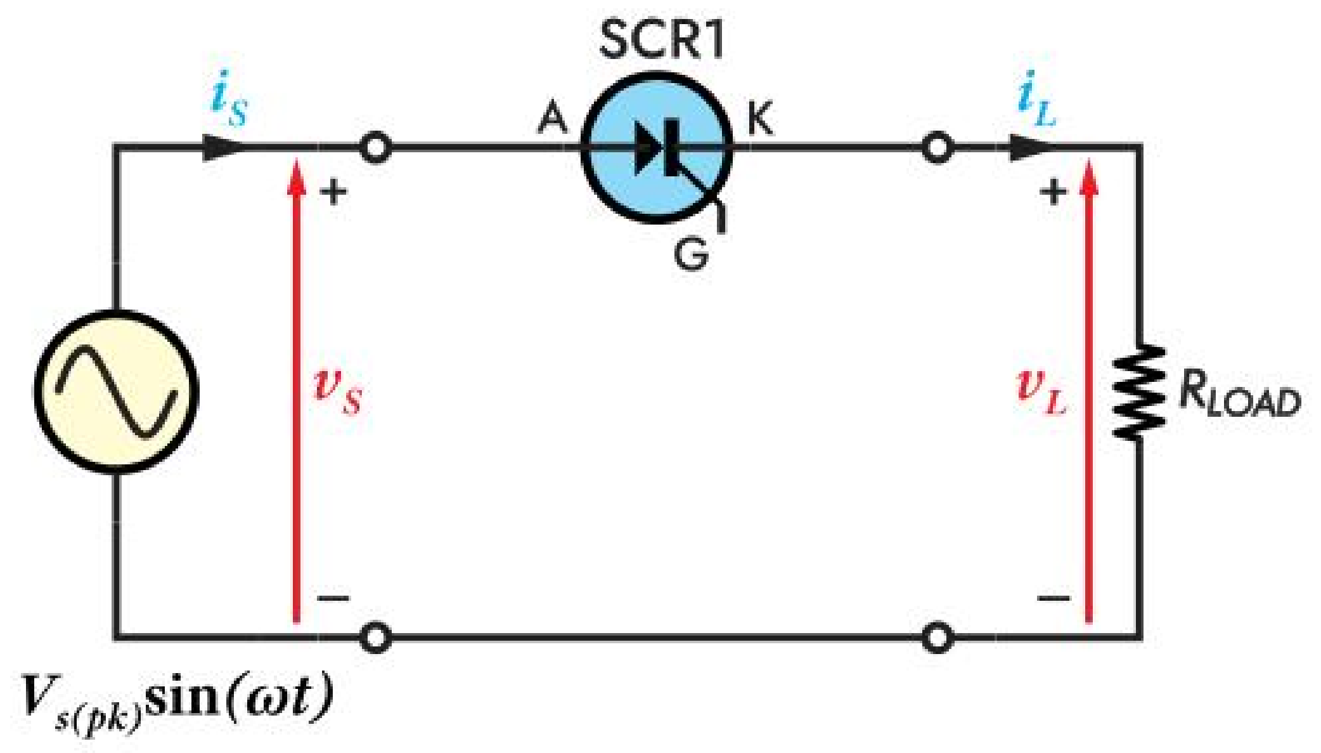
SCR1
i_S
A
K
G
i_L
+
v_S
−
+
v_L
−
R_{LOAD}
$V_{s(pk)}\sin(\omega t)$

Half-Wave Phase-Controlled Rectifier

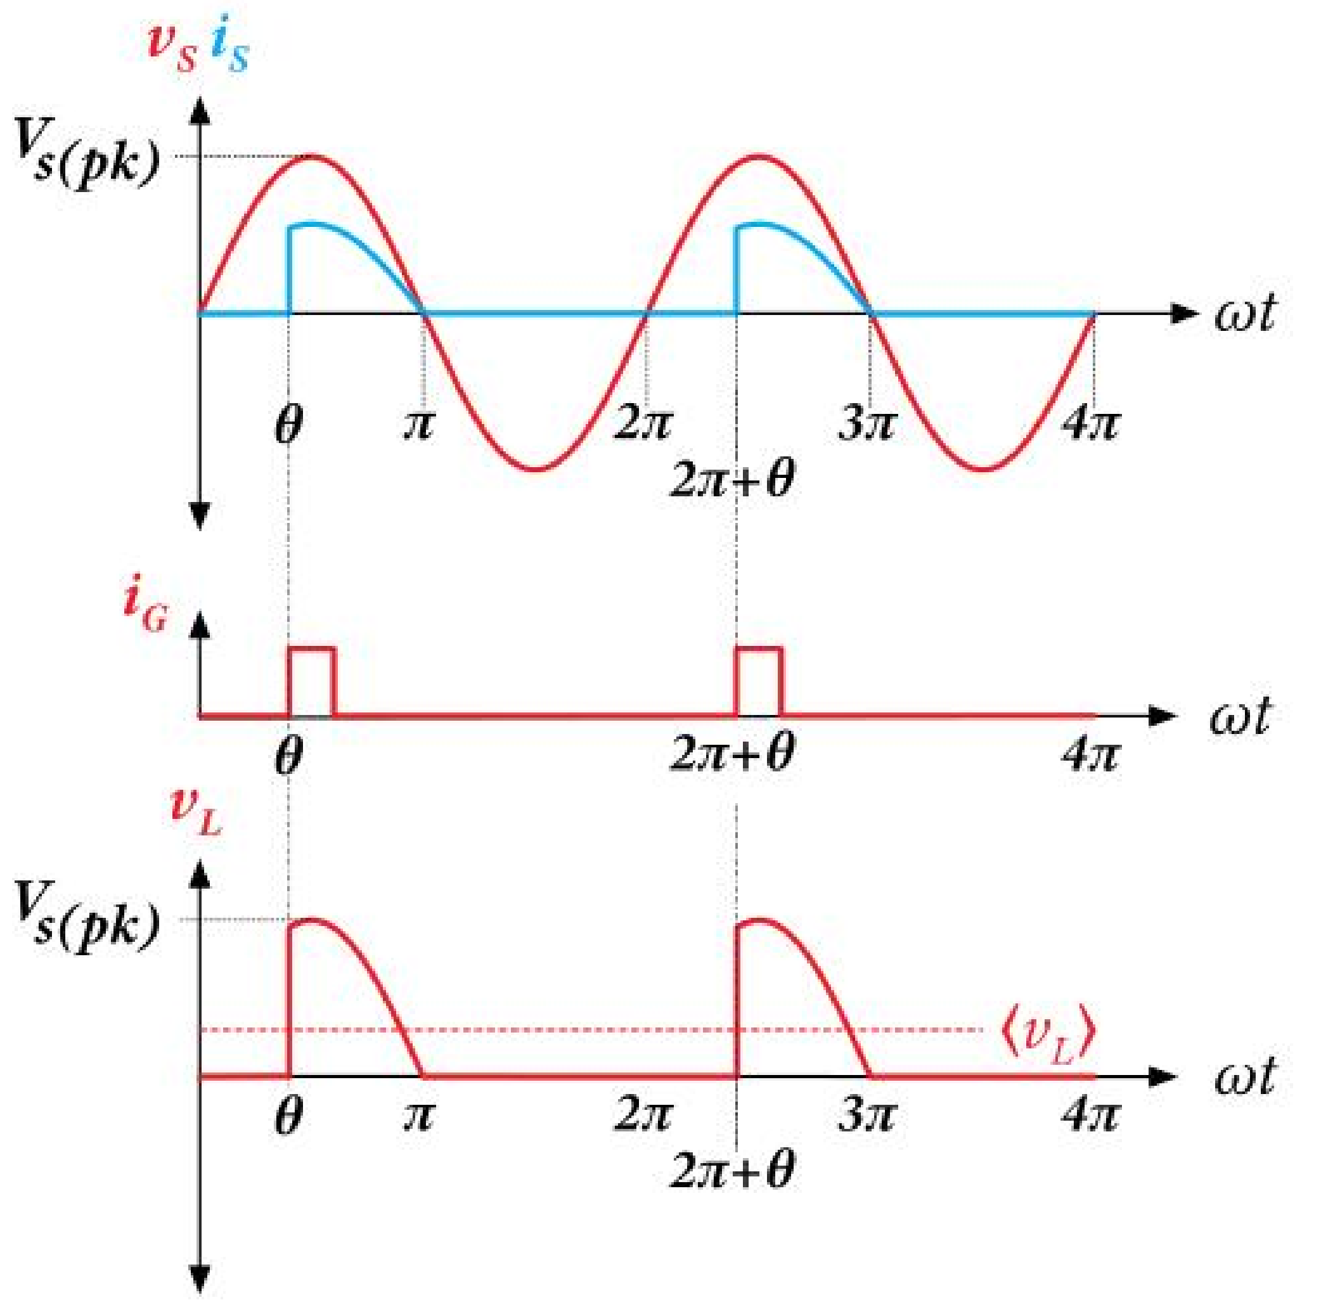

Fig.10: using a thyristor in place of the diode allows the output voltage of this half-wave rectifier to be controlled.

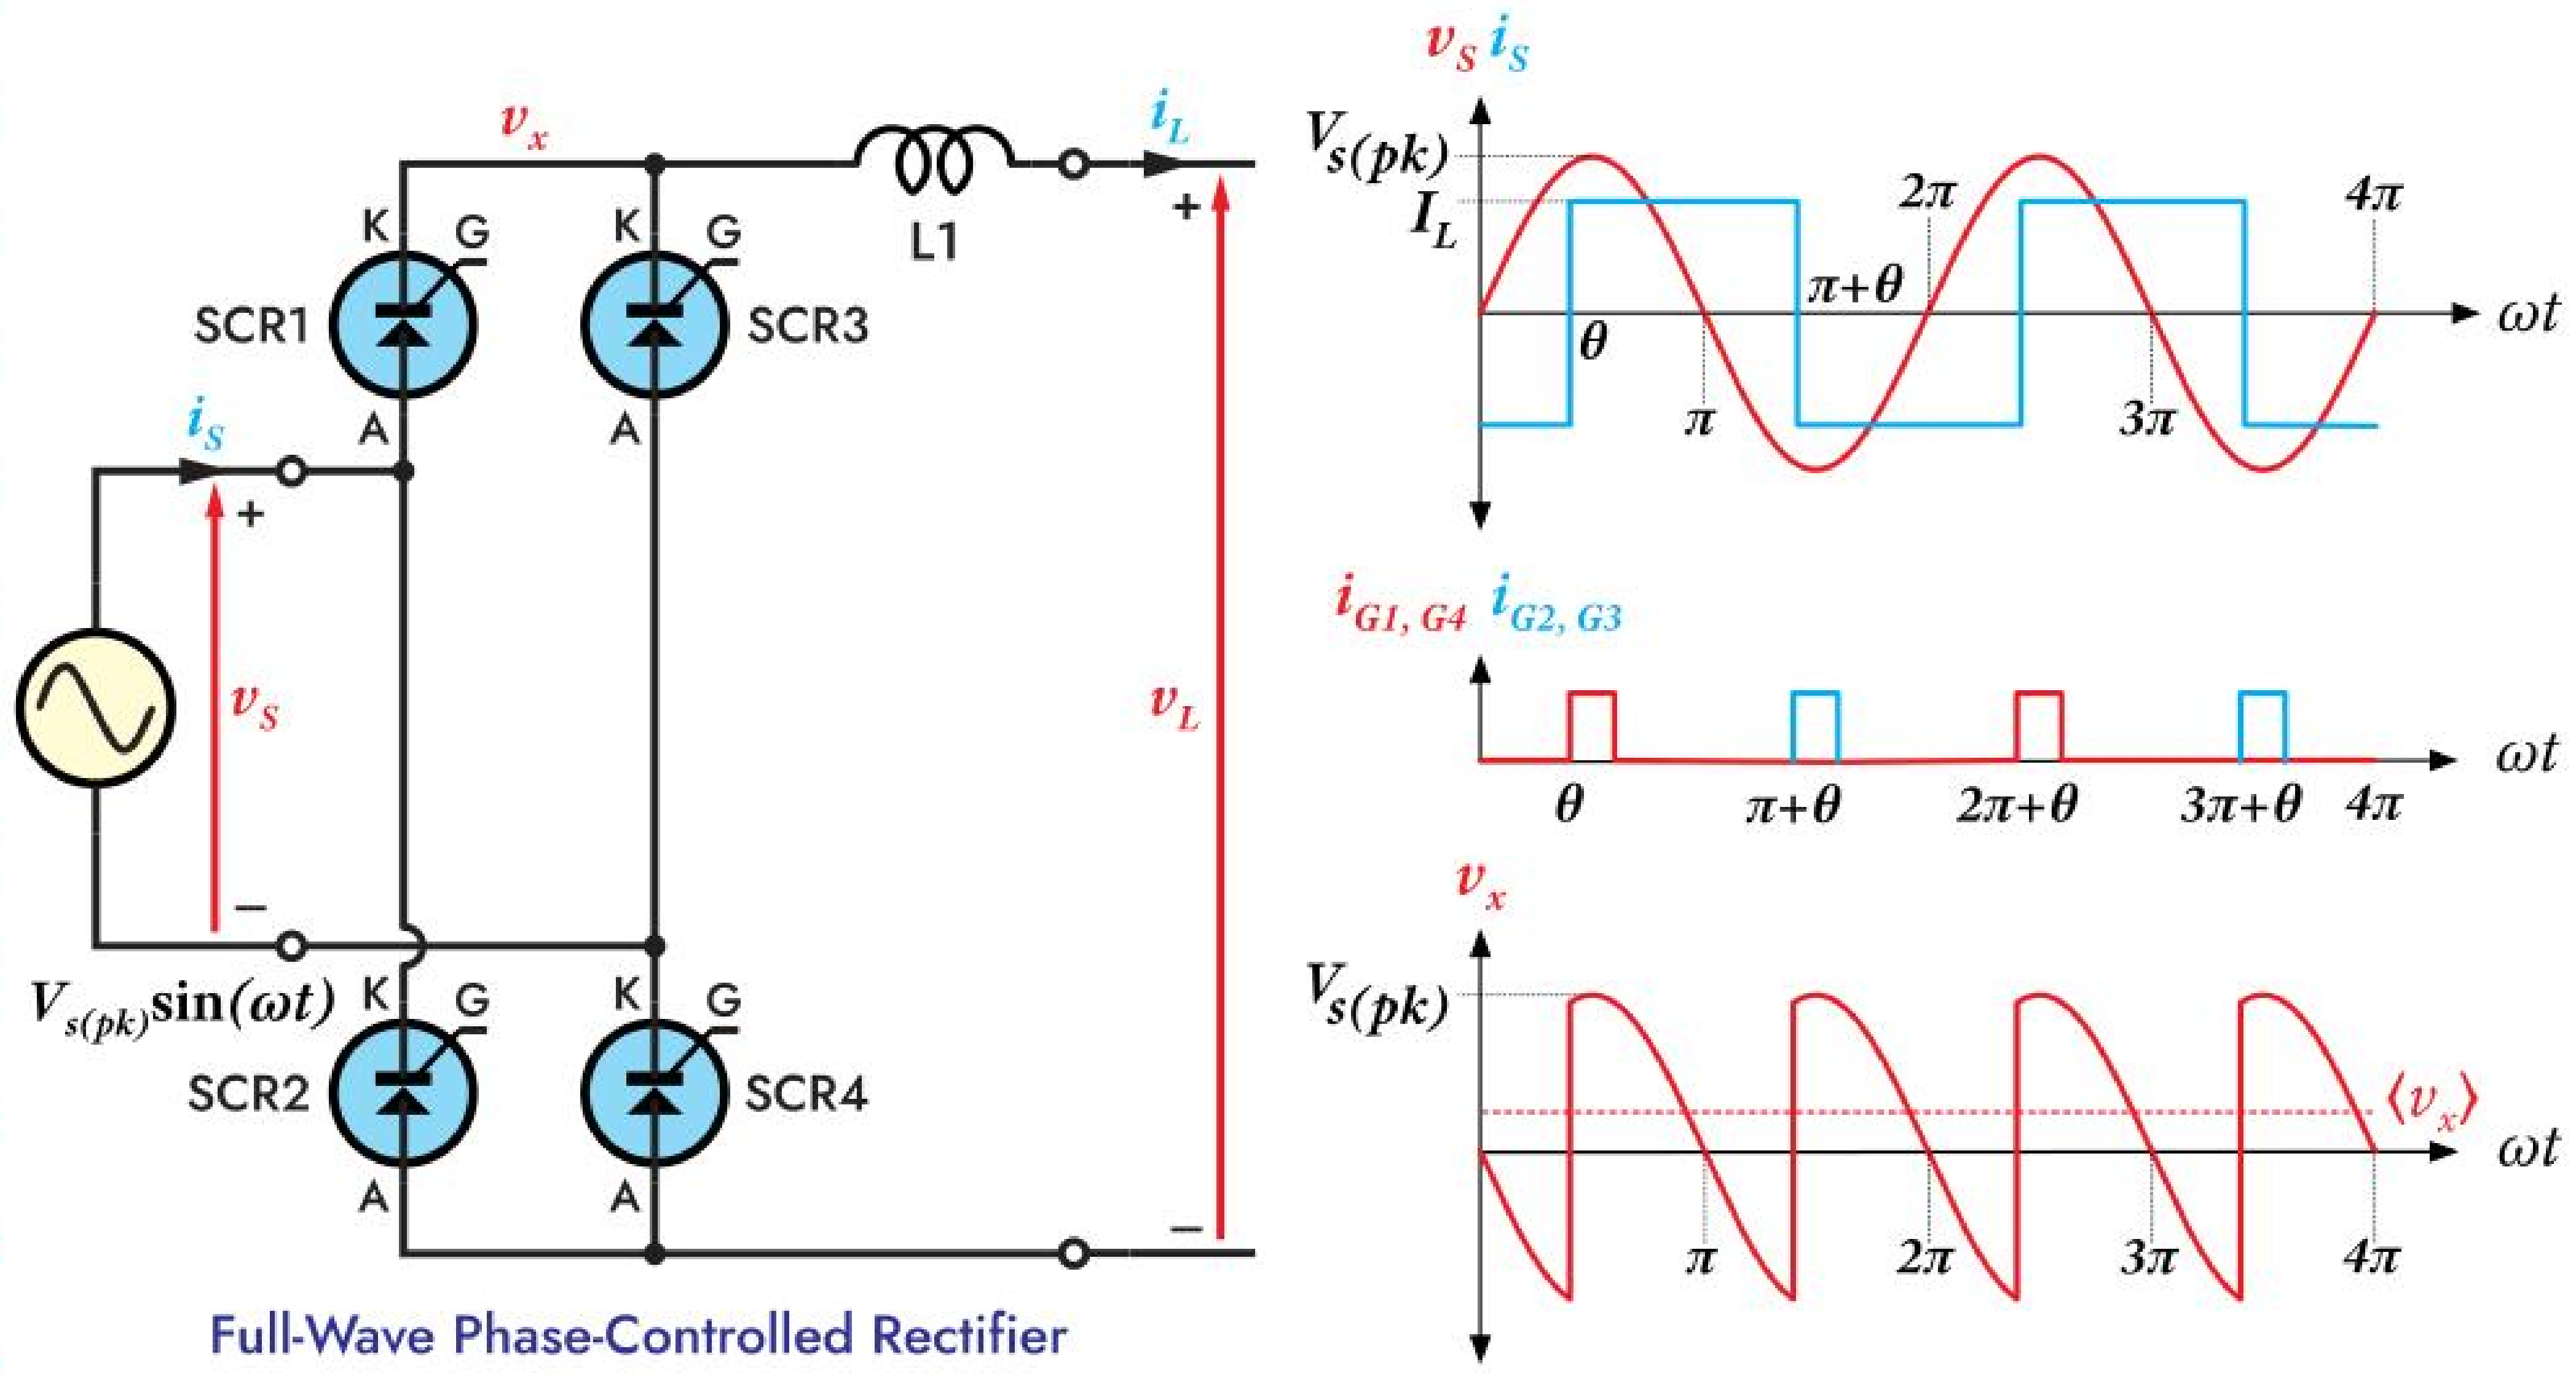
v_x
i_L
L1
+
K
G
SCR1
A
K
G
SCR3
A
i_S
+
v_S
−
v_L
$V_{s(pk)}\sin(\omega t)$
K
G
SCR2
A
K
G
SCR4
A
−
v_S i_S
$V_{s(pk)}$
I_L
2π
4π
$\pi+\theta$
θ
π
3π
ωt
$i_{G1, G4}$ $i_{G2, G3}$
ωt
θ
$\pi+\theta$
$2\pi+\theta$
$3\pi+\theta$
4π
v_x
$V_{s(pk)}$
$\langle v_x \rangle$
ωt
π
2π
3π
4π

Full-Wave Phase-Controlled Rectifier

Fig.11: the full-wave phase-controlled rectifier is capable of inversion, where energy is transferred from the load to the source.

The classical phase-control switch is the thyristor (sometimes called the silicon controlled rectifier, or SCR). You can think of a thyristor as a diode that will not conduct in the forward direction until the appropriate gate signal is applied.

When the gate is positively biased with respect to the cathode while the anode-cathode voltage is positive, the thyristor switches on and remains on, even if the gate bias is removed, until the current drops to zero. In this sense, it is a latching device.

In fact, the anode current must drop to zero for a short time (t_q) for the thyristor to recover its forward blocking capability. This time is in the order of tens to hundreds of microseconds, limiting the application of thyristors to fairly low-frequency applications.

Thyristors are very robust devices and are available in voltage ratings up to 6kV and current ratings up to 4kA. They have very good overload performance and, unlike most semiconductor switches, can be protected by fast fuses.

A more modest example, the 800V TN1605H-8G, is rated for a current of 16A RMS (8A average) and can withstand a non-repetitive half-cycle (10ms) surge of 160A.

Just as a matter of interest, this thyristor requires a gate current of 1.5mA to switch on (you would usually drive it at around 5mA to be sure), and has a t_q of 25μs at 25°C, rising to 85μs at 150°C.

Fig.10 shows a half-wave phase-controlled thyristor rectifier. In this case, the thyristor is switched on at a phase angle (sometimes called firing angle or delay angle) of θ. You can see that if θ is zero, the output will be identical to the half-wave rectifier (ie, $\langle v_L \rangle = V_{s(pk)} \div \pi$), but as the phase angle increases, the output voltage drops until it is zero when θ = π.

The relationship between phase angle and average voltage is not linear, due to the truncated sinusoidal shape of the voltage waveform. It can be shown that the output voltage is $\langle v_L \rangle = (V_{s(pk)} \div 2\pi) \times (1 + \cos[\theta])$.

Full-wave phase-controlled rectifier

The full-wave phase-controlled rectifier (Fig.11) has some very interesting properties, as we shall see. The thyristors are gated on at phase angle θ as before, with SCR1 and SCR4 on in the positive half cycle and SCR2 and SCR3 on in the negative one. Let's have a close look at what happens over a couple of cycles.

As we come to the end of a positive half-cycle (at phase angle π, let's say), SCR1 and SCR4 are conducting, but SCR2 and SCR3 have not yet been gated on to take over the constant current through the inductor.

This means the current continues to flow through SCR1 and SCR4, so they must remain conducting past phase angle π, and the voltage at v_x follows the input voltage negative.

When SCR2 and SCR3 switch on at phase angle π + θ, the current

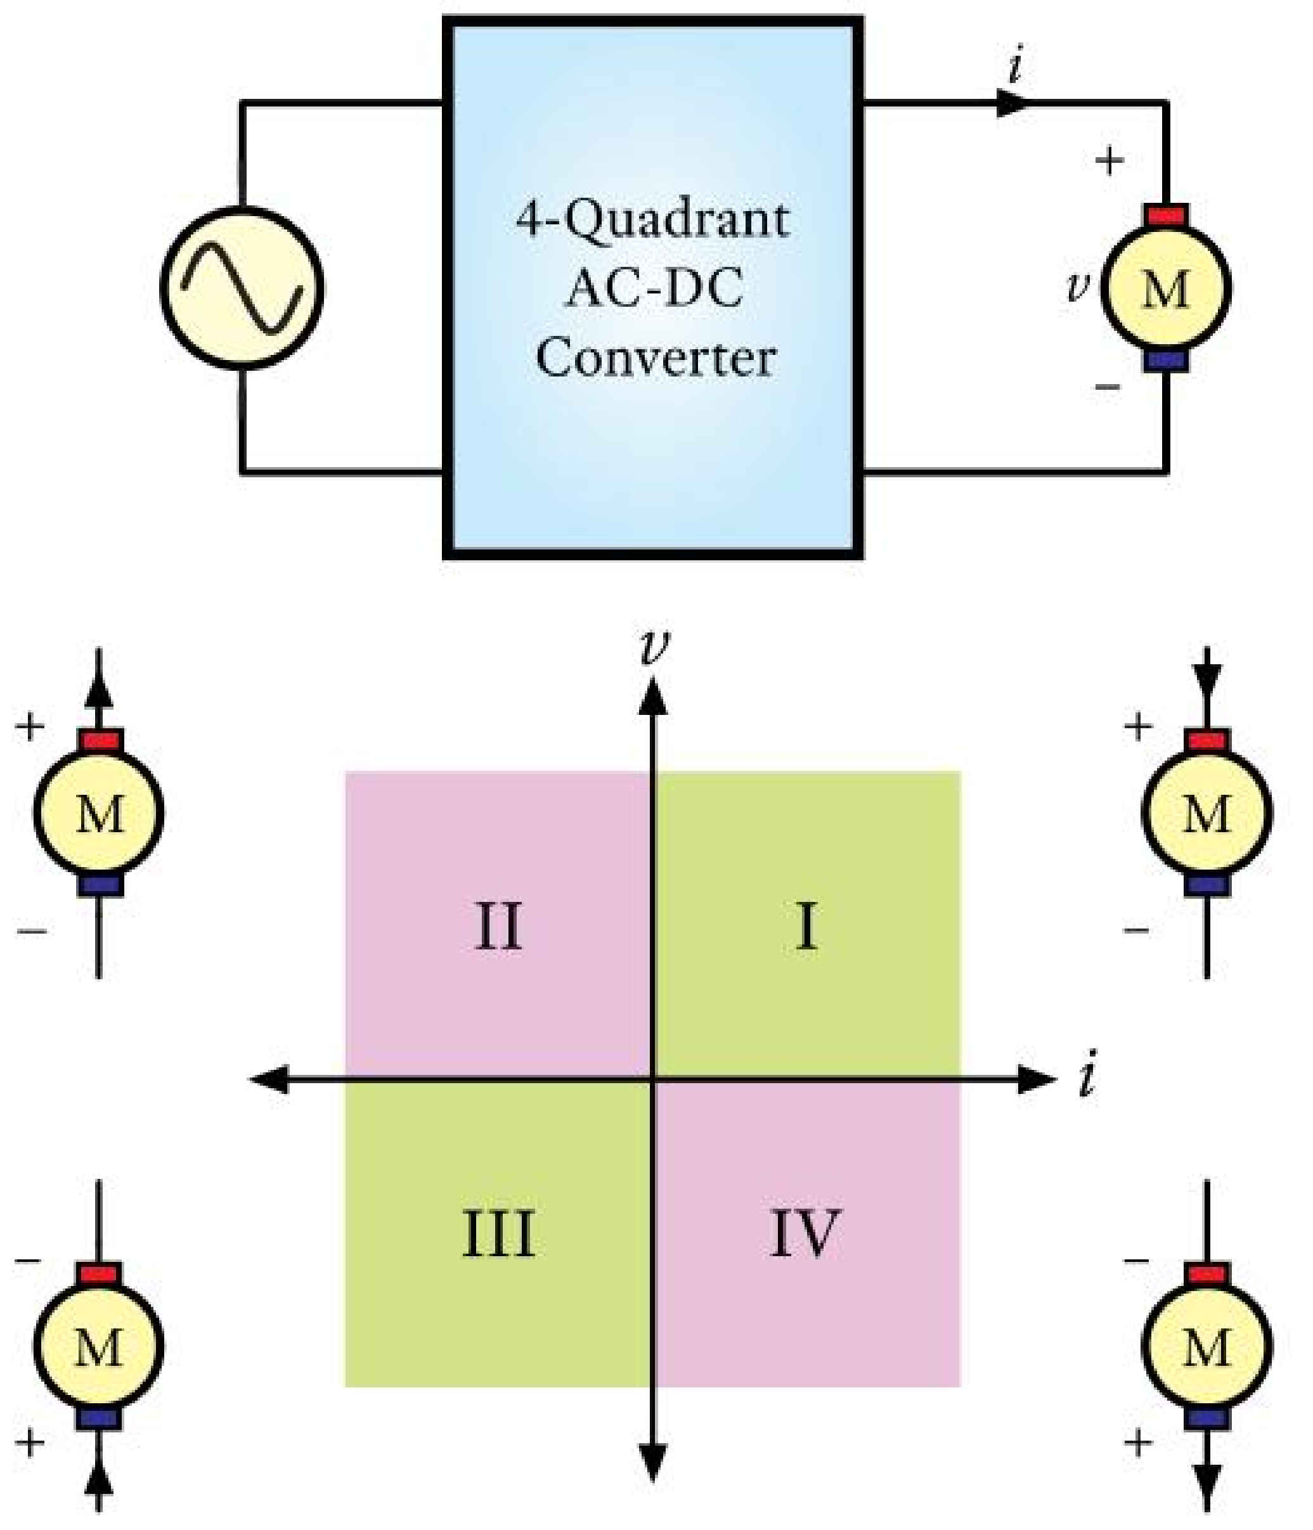

Fig.12: an AC-DC converter can, in theory, operate in any one of these quadrants. Inversion occurs in quadrants II and IV.

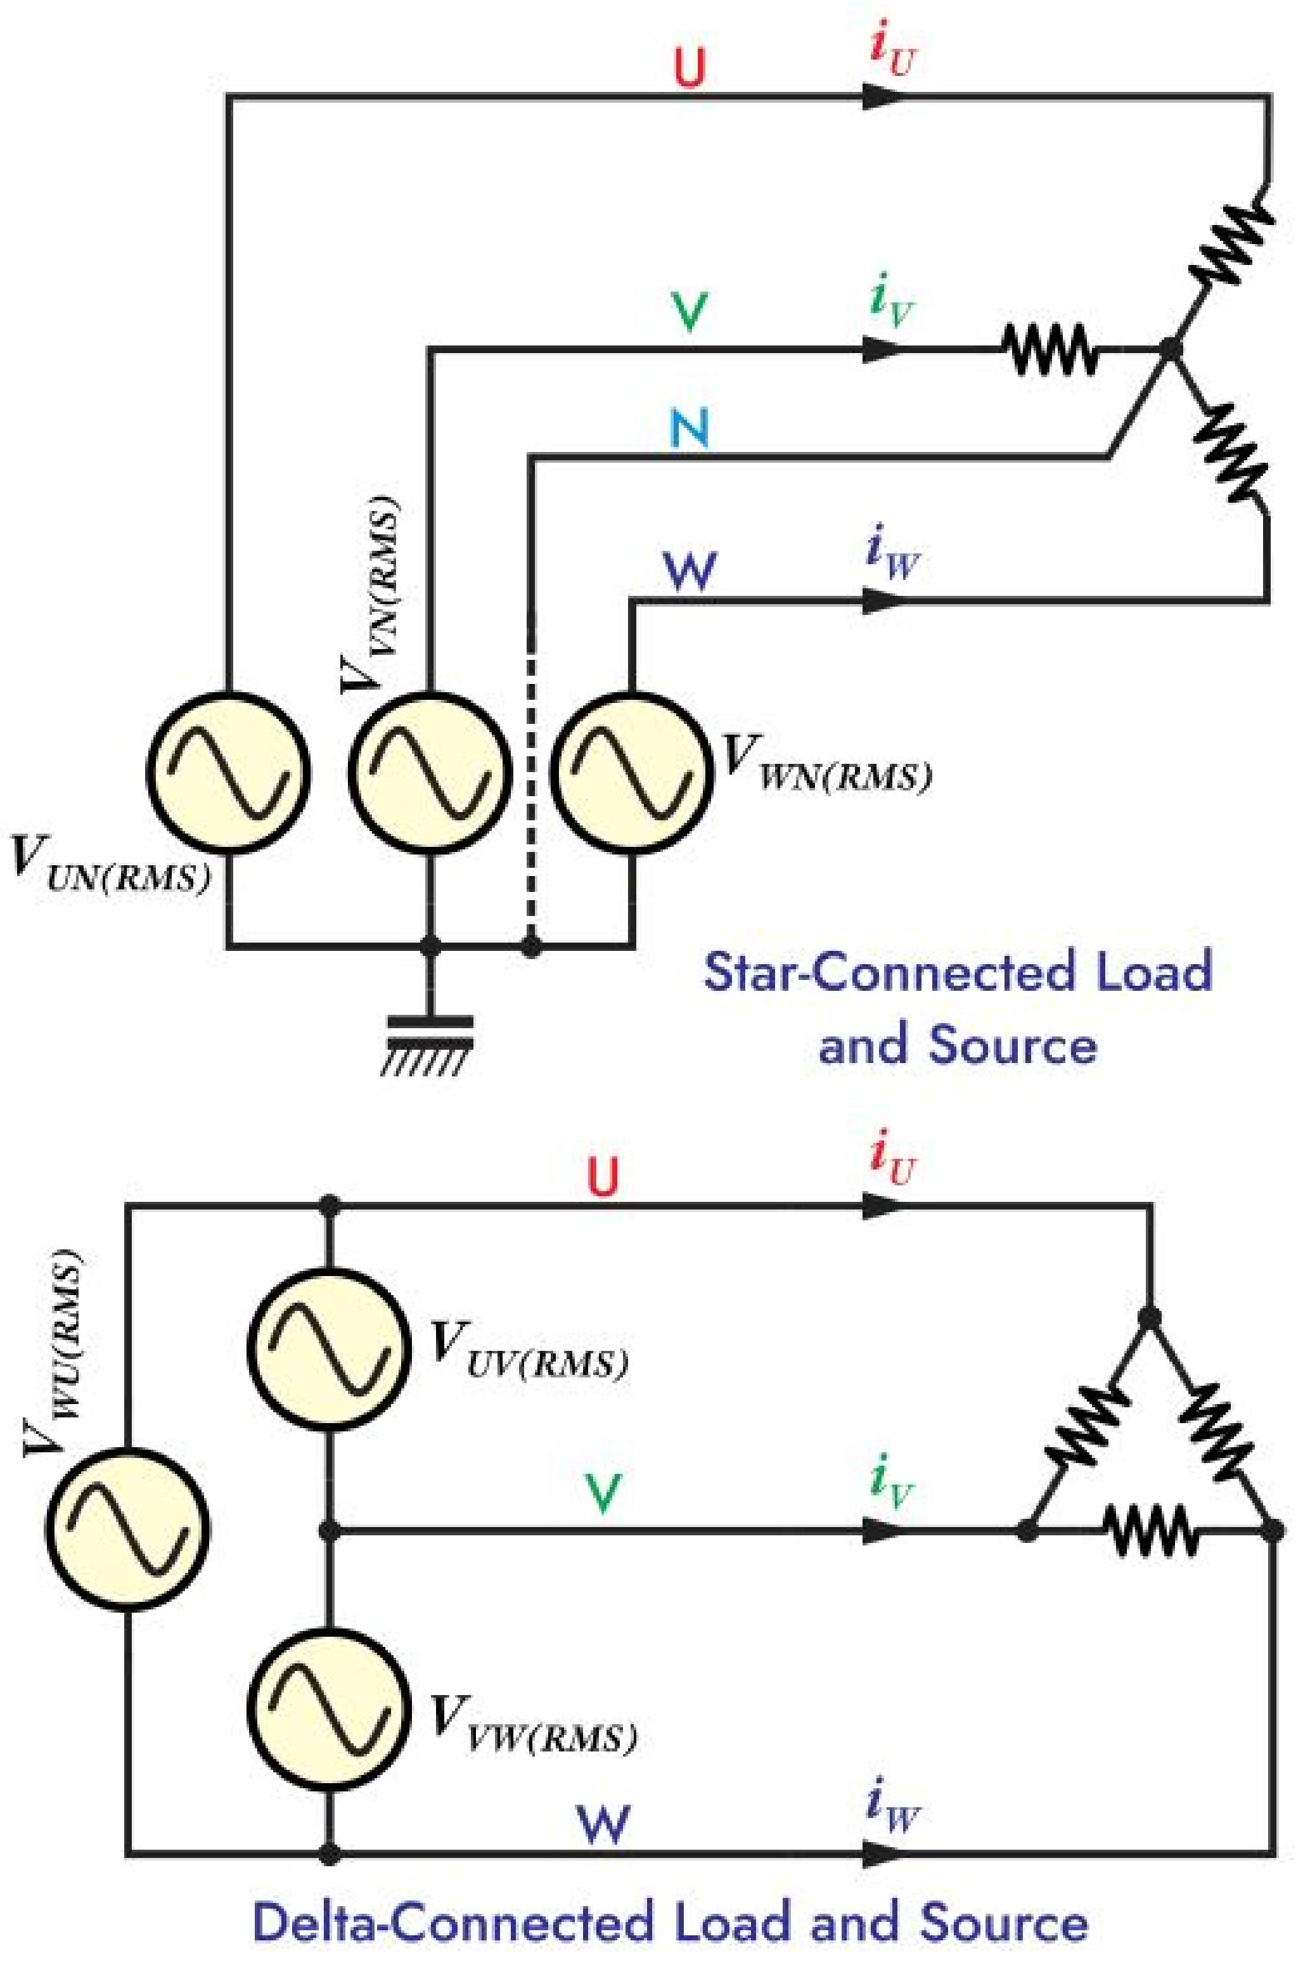

Fig.13: a three-phase source or load can be configured in star (Y) or delta (Δ), since no current flows in the Neutral wire if the phases are balanced.

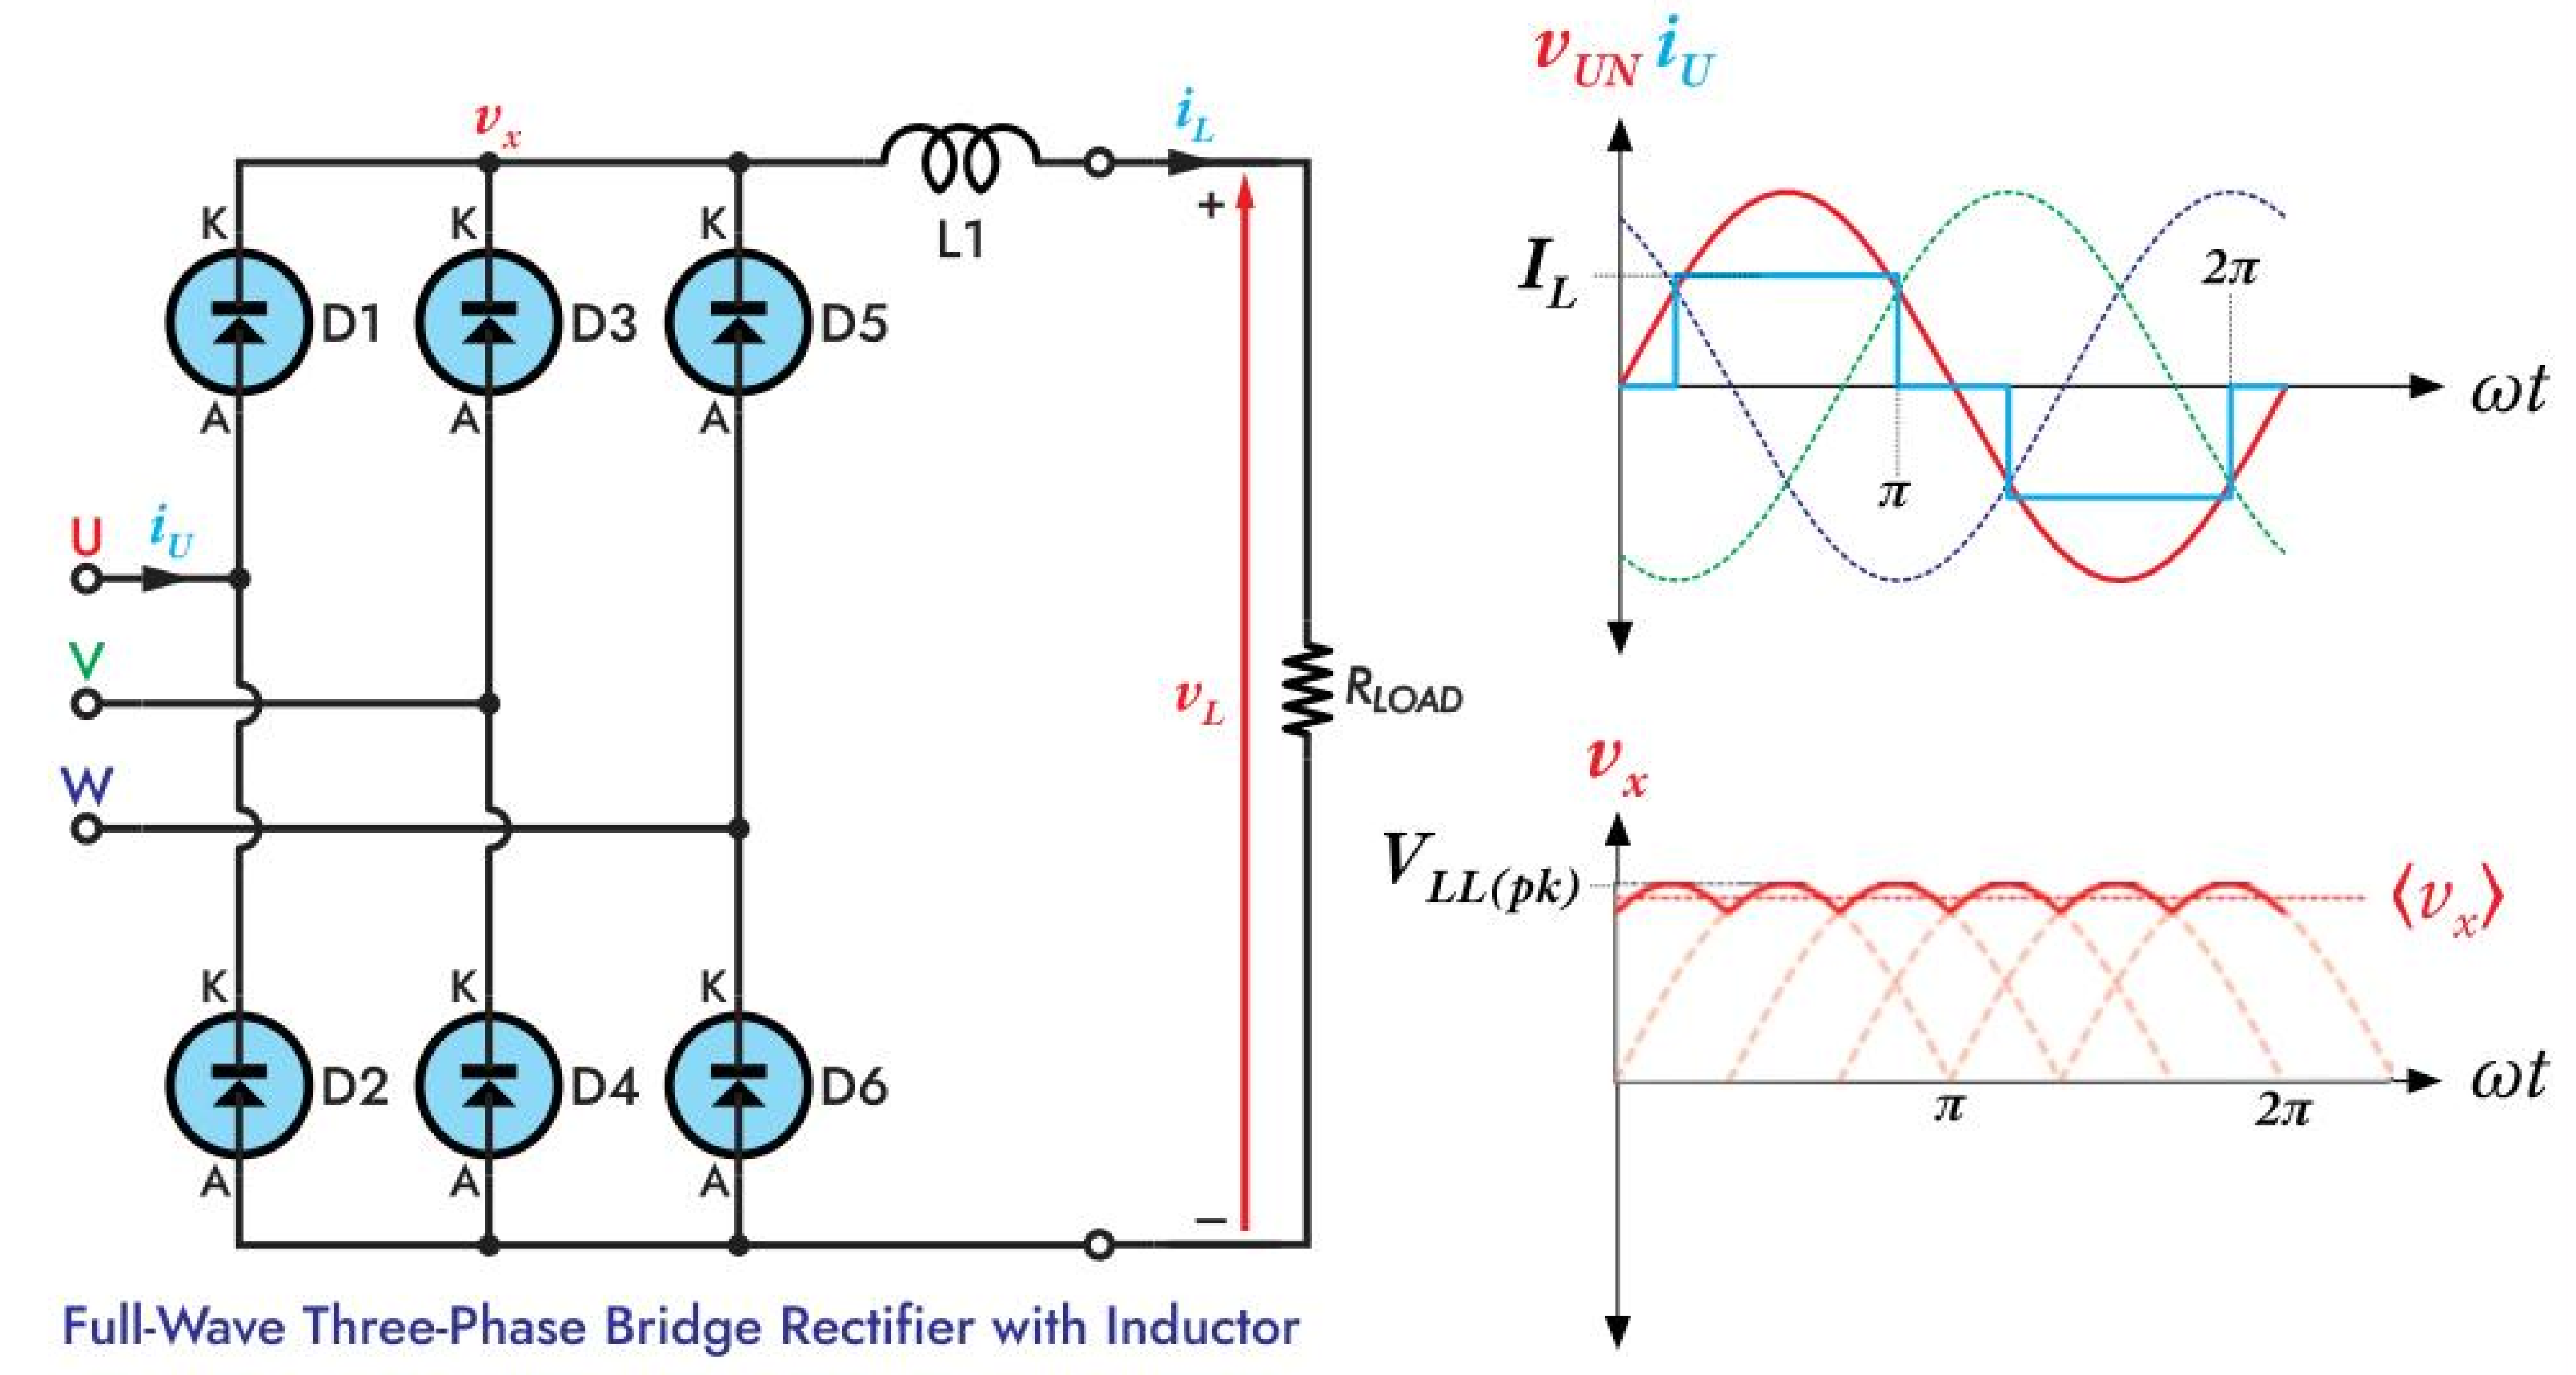

Fig.14: a three-phase rectifier has excellent low output ripple and a very good power factor, thanks to energy being delivered to the load six times per cycle.

commutates from SCR1 and SCR4 to SCR2 and SCR3, so the voltage flips positive abruptly. The same thing happens as the negative half-cycle comes to an end at phase angle 2π, although this time, it is SCR2 and SCR3 that remain on.

The upshot of this is that the average voltage, $\langle v_x \rangle$, and hence the average load voltage, can go negative. The input current is a square wave with a 50% duty cycle, as it was for the uncontrolled bridge rectifier in Fig.5, but now its phase can shift from being in phase with the input voltage when $\theta = 0$ to 180° out of phase when $\theta = \pi$.

Let's pause and take in what this means. On the load side, a positive DC current with a negative average voltage means negative power is 'dissipated' in the load! The same is true on the source side; at phase angles greater than $\pi/2$, the average input power is also negative, so power is delivered to the source.

This means that for phase angles greater than $\pi/2$, this circuit will transfer energy from the load to the source. This is known as inversion, and it can be very useful in practice.

All four quadrants

I will give a couple of examples where this can be useful, but first it is worth looking at the general case illustrated in Fig.12. Here, I have shown a generic 'four-quadrant' AC-DC converter powering a DC motor. The output voltage and current of this converter can each be positive or negative. This gives four possibilities, illustrated by the four quadrants in the V/I chart.

The quadrants are denoted by Roman numerals counter-clockwise from the top right. In quadrants I and III (shaded green), both the voltage and the current have the same sign, so the power is positive. In quadrants II and IV, the voltage and the current are of opposite polarities, so the power is negative.

I have positioned the current arrows on the motors so that they are always on the most positive terminal, but they are consistent with the upper diagram in that positive current flows down and negative current flows up.

You can see that in quadrants I and III, the current flows into the most positive terminal, so the motor will be consuming power and driving the load (possibly in opposite directions, depending on the motor type).

In quadrants II and IV, the current emerges from the most positive terminal of the motor, so the motor must be behaving as a generator and exporting power.

Getting back to the full-wave phase-controlled rectifier in Fig.11, we can see that this operates in two quadrants (I and IV, because the current is always positive). If the load on this converter was a DC motor with a high-inertia mechanical load, quadrant I could be used to drive the load, and quadrant IV could provide regenerative braking.

The reversed voltage applied to the motor creates a retarding torque that brings the motor and load to a stop quickly, since the energy stored in the rotating mass is transferred to the source. The motor will come to a stop much faster than it would if left to coast, losing its stored energy only to friction and windage.

I have also seen this circuit used to quickly switch off (and therefore de-magnetise) a large electromagnet, moving the energy stored in the magnet's inductance to the supply much faster than it would if we relied on the freewheeling effect of a full bridge rectifier. This allowed the electromagnet, which was picking up and moving steel components, to release them promptly on command.

We should note here that the inversion described here is not sufficient to create a DC-AC converter. The AC source must be present for this circuit to work, and power can only flow back to the source while it is present. I will cover DC-AC inverters in a future article.

Three-phase systems

The electromagnet example I mentioned above was actually fed from a three-phase supply. I won't go into multi-phase rectifiers in any great detail, since they are really only used in industrial applications, but I will touch on them for completeness.

A three-phase voltage source is just three sinusoidal voltages with equal amplitude and frequency, but shifted in phase by one third of a cycle.

I have drawn these three sources and loads (at the top of Fig.13) in a slightly unconventional manner, but you can probably see why this arrangement is called a star or Y configuration. The centre point of the star is the Neutral connection. The phase-to-Neutral voltages are shifted from each other by $2\pi/3$ radians (or 120°), and this brings some very useful benefits.

Firstly, if the load in each phase is balanced, the sum of the three line currents is zero, and no current flows in the Neutral wire, which is why I have shown it dotted. In fact, many three-phase loads such as motors don't even have a Neutral terminal.

On top of this, for balanced loads, any current harmonics that are multiples of three (the 3rd, 6th, 9th etc) also cancel to zero so, they don't contribute to apparent power, meaning you get a power factor advantage for free.

Given that the Neutral is unnecessary if the load is balanced, we could think about line-to-line voltages rather than line-to-Neutral voltages, and redraw the circuit as shown at the bottom of Fig.13. Again, I have drawn it unconventionally, but this configuration is called a delta (Δ) arrangement because the elements are arranged in a triangle.

You can have a star-connected load with a delta-connected source and vice versa.

The line-to-line (or phase-to-phase) voltages are the sum of the two relevant line-to-Neutral voltages. I will leave the maths out, but it is easy enough to show that the line-to-line voltages are larger than the line-to-Neutral voltages by a factor of $\sqrt{3}$, and displaced from them in phase by $\pi/6$ radians (30°).

The nominal line-line voltage for domestic three-phase supplies in the UK is therefore 400V, corresponding to $\sqrt{3}$ times the 230V nominal phase-to-Neutral voltage.

Before you rush to correct me, I am well aware the typical line-to-line voltage is closer to 415V in many areas of the country, as we continue to transition from the old 240V/415V standard to the newer 230V/400V standard. In any case, 240V/415V is within the allowable range of the 230V/400V standard and vice versa.

A three-phase full-wave rectifier

Fig.14 shows a three-phase full-wave rectifier with an inductor. The upper graph shows the U-phase line-Neutral input voltage and the U-phase line current, with the V-and W-phase line-Neutral voltages shown dotted. The lower graph shows the voltage at v_x, with the six half-cycles that contribute to it shown dotted.

The average voltage at v_x and hence the average load voltage $\langle v_L \rangle = 3v_{LL(pk)} \div \pi$ or approximately 0.96 times the peak line-line voltage. The load voltage ripple is obviously much lower than for the single-phase rectifier, since there are now six 'pulses' of voltage each cycle instead of two.

The power factor for the three-phase rectifier is also better than the single-phase case, as you might expect by observing that the line current waveform looks more sinusoidal with its 'stepped' shape. The power factor also turns out to be equal to $3/\pi$ (0.96) for this topology, so quite close to unity.

You can also use thyristors in place of the diodes to create a phase-controlled version of this rectifier. It behaves in much the same way as its single-phase counterparts with phase angles above $\pi/2$, producing negative output voltages. Like the single-phase case, it can operate in quadrants I and IV.

That's all for this month. Next month, we will continue to look at AC-DC converters, with a focus on power factor correction. While we are on the topic of being responsible with regard to the power grid, we will also touch on the basics of electromagnetic interference (EMI) filtering. PE

PoLabs

For more information or software download please visit

www.poscope.com/epe

PoKeys
Connect, Control

- USB
- Ethernet
- Web server
- Modbus
- CNC (Mach3/4)
- IO
- PWM
- Encoders
- LCD
- Analog inputs
- Compact PLC

Stepper motor drivers
Drive

- up to 256 microsteps
- 50 V / 6 A
- USB configuration
- Isolated
- up to 32 microsteps
- 30 V / 2.5 A

PoScope Mega1+
PoScope Mega50
Measure

- up to 50MS/s
- resolution up to 12bit
- Lowest power consumption
- Smallest and lightest
- 7 in 1: Oscilloscope, FFT, X/Y, Recorder, Logic Analyzer, Protocol decoder, Signal generator

The Fox Report

Barry Fox's technology column

Balcony solar

For years, I have been banging on about simple, low-cost solar. One or two panels stood or hung in the garden or over a balcony, plugged into an inverter that converts 20V DC to 230V AC and pushes it into the house mains via an existing standard mains socket.

No official permissions or paperwork and a free few hundred watts that – while the sun shines – cover the house 'base load' continually drawn by a fridge, computer or TV on standby.

Of course, it isn't that simple. The inverter has to synchronise its frequency with the mains and not upset any circuit breakers. If the inverter is feeding 230V through a three-pin mains plug, it has to disconnect within milliseconds of unplugging from the mains socket to avoid becoming a shock hazard.

Technically, this is no longer difficult. 'Balcony solar' is now big in Germany and other continental countries, where safe plug-and-push power kits are on sale. Hang the panel(s) over a balcony, plug a couple of DC leads into a weatherproof (or internal) inverter and run a mains lead with a standard mains plug into a standard wall socket.

Chinese factories are churning out push-power kits, which retailers are badging, and house-owners or renters are cash-and-carrying from a supermarket. When you move home, you just unplug and go.

Better late than never, the UK government is now moving to allow similar sales in Britain. The official announcement (**https://pemag.au/link/acbz**), dated 24th of March, is long on hyperbole and enthusiastic quotes from Uncle Tom Cobley and all, but short on hard facts. I quote:

'Plug-in' solar panels to be in shops within months, offering households a chance to significantly cut energy bills.

More households will be helped to save money on bills through plug-in solar panels as the government vows to go further and faster on clean energy…

… Germany is seeing around half a million new devices plugged in per year.

The free solar power can be used directly through a mains socket like any other device, without an installation cost, thereby reducing the amount of electricity taken from the grid and cutting energy bills.

The easy-to-install technology could save many households significant amounts on their energy bills and help make the UK less reliant on global fossil fuel markets.

Georgina Hall, Corporate Affairs Director at Lidl GB, is quoted as saying:

At Lidl GB, we are committed to making sustainable living affordable for everyone and we welcome the government's move to modernise regulations in the UK. Updating the regulatory landscape for this 'plug-and-play' technology is a positive step towards empowering British households to manage their energy costs and support the nation's net-zero ambitions.

As someone who bought and struggled to set up one of Lidl's robot lawn mowers (which failed with unhelpful error messages after one season), I must admit to having some reservations here.

The government says it "will work with the Energy Networks Association, DNOs and Ofgem to update the G98 distribution code and wiring regulations BS 7671 to allow UK households to connect <800W plug-in solar panels to domestic mains sockets, without the need for an electrician and with tailored safety standards".

Fine words, but I suspect the government may have underestimated the mass of paperwork and official permissioning that need to be sorted. There has been – and rightly so – a jungle of restrictions on who can connect what to the national grid via domestic wiring, fusing and metering.

It's a fact of electronics that if someone can find a way to connect something incorrectly, they will. Will the 800W power ceiling apply to individual plugs or whole houses? Can separate inverters be fed into separate sockets?

High-tech equipment is now sold like sweets. I recently stood behind someone in Argos who was hoping to get his crashed or infected laptop working again under guarantee.

Although it is not a safety issue, how to correctly wire the low-voltage plugs and sockets from the panels can be confusing, with the negative and positive, male and female, outer casings and internal pins; it's easy to get it wrong.

The plugs won't grip thin wire cores. So the kits sold will need to come with all plugs and wires pre-fitted and weatherproofed.

Mains plugs will have to be pre-fitted to avoid incorrect Live, Neutral, Earth wiring, which could play havoc with safety circuit breakers.

From equipment already listed by Amazon, the likely starting price will be around £500 for two panels and a microinverter. As an immediate, impatient experiment, I bought an inverter direct from China for a little under £50 and two 100W panels for £50 each, plus a few metres of two-core garden mains cable and a selection of connectors. The total cost was around £175.

My inverter came with Chinglish instructions and a very confusing description of what its LED is trying to tell me:

- Long bright green light: micro inverter to normal
- Green light flashing: micro inverter power
- Long bright red light: an island protection and fault

Red light flashing:

1. AC voltage under-voltage protection (flash a red light)
2. AC voltage overpressure protection (red light flashing 2)
3. Frequency protection (red light flashes three times)

Very likely, some badged Brit kits will come with instructions that blindly copy this or similarly garbled text.

After some trial and error, I think the basics are: when the inverter LED shows steady green, it is correctly connected to a working panel and live mains; when it flashes green at regular intervals, it is receiving

The Fox Report

Barry Fox's technology column

good DC and pushing AC into the mains; when the LED turns steady red, something is wrong.

My microinverter is, on paper, rated too low in wattage for the two 100W panels. In practice, it seems well able to cope, even on sunny days.

To see if the panel is pushing useful power into the mains when the sun is shining and the inverter is flashing green, I checked the house smart meter while using a switch plug to connect and disconnect the inverter from the mains.

It takes a while for the meter to register the change (by which time a fridge may have cut in or out), but with the sun shining, the two panels seem to be cutting consumption by around 2p per hour, so around 1p per panel. With, let's hope, an average of 8 hours sun or bright light per day, that's around £30 a year per panel, so £60 for a two-panel kit.

Perhaps commercial kits will do better, and perhaps the price of a two-panel kit will eventually fall to £199. Even so, I suspect that many people using balcony solar will earn less than they expected.

Personally, though, I am happy just to see the LED flashing green and telling me I am getting at least some free power from the sun. PE

Source: RobbieIanMorrison (https://commons.wikimedia.org/wiki/File:Balcony_solar_panels_in_Wesenberg,_Germany_a.jpg)

ETI BUNDLE (1) Teach-In 3, 4 and 5 – all on CD-ROM – only £18.95

ELECTRONICS TEACH-IN 3 – CD-ROM
Mike & Richard Tooley

The three sections of the *Teach-In 3* CD-ROM cover a huge range of subjects that will interest everyone involved in electronics – from newcomers to the hobby and students to experienced constructors and professionals.

The first section (80 pages) is dedicated to *Circuit Surgery*, *EPE/PE's* regular clinic dealing with readers' queries on circuit design problems – from voltage regulation to using SPICE circuit simulation software.

The second section – *Practically Speaking* – covers hands-on aspects of electronics construction. Again, a whole range of subjects, from soldering to avoiding problems with static electricity and identifying components is covered. Finally, our collection of *Ingenuity Unlimited* circuits provides over 40 circuit designs submitted by readers.

The CD-ROM also contains the complete *Electronics Teach-In 1* book, which provides a broad-based introduction to electronics in PDF form, plus interactive quizzes to test your knowledge and TINA circuit simulation software (a limited version – plus a specially written TINA Tutorial).

The *Teach-In 1* series covers everything from electric current through to microprocessors and microcontrollers, and each part includes demonstration circuits to build on breadboards or to simulate on your PC.

ELECTRONICS TEACH-IN 4 – CD-ROM
A BROAD-BASED INTRODUCTION TO ELECTRONICS
Mike & Richard Tooley

The *Teach-In 4* CD-ROM covers three of the most important electronics units that are currently studied in many schools and colleges. These include, Edexcel BTEC level 2 awards and the electronics units of the Diploma in Engineering, Level 2.

The CD-ROM also contains the full *Modern Electronics Manual*, worth £29.95. **The Manual** contains over 800 pages of electronics theory, projects, data, assembly instructions and web links.

This will appeal to anyone interested in learning about electronics – hobbyists, students or professionals.

ELECTRONICS TEACH-IN 5 – CD-ROM
JUMP START
Mike & Richard Tooley

This CD-ROM includes 15 design and build circuit projects for newcomers or those following courses in school and colleges.

The projects are: ■ Moisture Detector ■ Quiz Machine ■ Battery Voltage Checker ■ Solar-Powered Charger ■ Versatile Theft Alarm ■ Spooky Circuits ■ Frost Alarm ■ Mini Christmas Lights ■ iPod Speaker ■ Logic Probe ■ DC Motor Controller ■ Egg Timer ■ Signal Injector Probe ■ Simple Radio Receiver ■ Temperature Alarm.

PLUS

PIC'n' Mix – starting out with the popular range of PIC microcontrollers and *Practically Speaking* – tips and techniques for project construction.

The CD-ROM also contains:
■ Complete *Teach-In 2* book, a practical introduction to PIC microprocessors
■ MikroElektronika, Microchip and L-Tek PoScope software.

Three Teach-ins for the great price of

£18.95

PLUS you also get the contents of the free CD-ROM from each issue...

... so that's another TWO Teach-Ins and The Full Modern Electronics Manual!

What a Bargain!!

Practical Electronics PCB SERVICE

PROJECT	CODE	PRICE
JULY 2026		
RP2350B Development Board (assembled)	07107251A	£24.99
SmartProbe PCB	P9054-04	£4.99
SWD Adaptor board for the SmartProbe	P9045-A	£2.99
Audio Out New All-Pass Filter PCB	AOJUL26-1	£4.99
Audio Out Alps Polarity Switch PCB	AOJUL26-2	£2.99
Audio Out E-switch Polarity Switch PCB	AOJUL26-3	£2.99
JUNE 2026		
RGB LED 'Analog' Clock PCB (black)	19101251	£7.99
USB Power Adaptor PCB (set of two)	18101251	£2.99
433MHz Receiver Module PCB	15103252	£2.99
Audio Out All-Pass Filter PCB set (three PCBs)	AO-JUN-26	£6.99
MAY 2026		
SSB Shortwave Receiver three PCB set	CSE250202-4	£15
Tool Safety Timer PCB	10104251	£3.99
433MHz Transmitter Module PCB	15103251	£2.99
APRIL 2026		
Versatile Battery Checker PCB	11104251	£3.99
Versatile Battery Checker front panel (black, 0.8mm)	11104252	£4.99
Power Rail Probe PCB	P9058-1-C	£3.99
Pico 2 Audio Analyser PCB/panel (black)	04107231	£3.99
MARCH 2026		
USB-Programmable Frequency Divider PCB (blue)	04108241	£3.99
Rotating Light PCB (black)	09101251	£2.99
433MHz Transmitter PCB	15103251	£2.99
Arduino Uno Game Cartridge PCB (Cool Beans)	CBmar26uno	£3.99
Arduino Nano Game Cartridge PCB (Cool Beans)	CBmar26nano	£2.99
Silent Stomp Switch PCB pair (Audio Out)	AO-MAR26-*	£3.99
FEBRUARY 2026		
Power LCR Meter PCB	04103251	£7.99
Versatile Waveform Generator PCB	04104251	£5.99
JANUARY 2026		
High-Bandwidth Differential Probe	9015-D	£5.99
NFC-Programmable IR Remote Keyfob	15109231	£2.99
Pre-driver Overdrive Pedal main PCB	AO-JAN26-1	£4.99
Updated jack PCB for the Pre-driver	AO-JAN26-2	£3.99
DECEMBER 2025		
Variable Speed Drive (black)	11111241	£11.99
Digital Capacitance Meter	04111241	£5.99
Battery-Powered Model Train transmitter	09110241	£3.99
Battery-Powered Model Train through-hole receiver	09110242	£3.99
Battery-Powered Model Train SMD receiver	09110243	£3.99
Battery-Powered Model Train battery charger	09110244	£3.99
Transmanium Fuzz Box jack PCB	AO-DEC-25	£3.99
NOVEMBER 2025		
Capacitor Discharger	9047-01	£4.99
5MHz, 40A Current Probe	9049-01	£5.99
Transmanium Fuzz Box main PCB	AO-NOV-25	£4.99
OCTOBER 2025		
Seaside Sound Simulator PCB (blue)	01111241	£7.99
Fender Jazz Bass pickguard	23109241	£5.99
J&D Luthiers T-Style pickguard	23109242	£5.99
Music Man Stingray pickguard	23109243	£5.99
Fender Telecaster pickguard	23109244	£4.99
3D Printer Filament Dryer control PCB	28110241	£4.99
SEPTEMBER 2025		
Automatic LQ Meter PCB	CSE240203A	£5.99
Automatic LQ Meter front panel (black)	CSE240204A	£5.99
Compact Hi-Fi Headphone Amplifier PCB (blue)	01103241	£4.99
Discrete Ideal Bridge Rectifier TH PCB	18108241	£2.99
Discrete Ideal Bridge Rectifier SMD PCB	18108242	£2.99
Dual Rail Supply Protector PCB	18109241	£3.99
Audio Out Cascode Amplifier PCB	CASCODEA	£4.99
Audio Out Cascode Amplifier Power Supply PCB	CASCODEP	£4.99
AUGUST 2025		
Micromite Explore-40 PCB	07106241	£2.99
Micromite Explore-40 Audio Breakout PCB	07101222	£2.99
8-Channel Learning IR Remote Receiver PCB	15108241	£5.99
FlexiDice PCB set	08107241-2	£4.99
JULY 2025		
Compact OLED Clock & Timer PCB	19101231	£4.99
180-230V DC Motor Speed Controller PCB	11104241	£9.99
Mains Power-Up Sequencer PCB (same as Feb '25)	10108231	£11.99
JUNE 2025		
Styloclone (case version)	23106241	£6.99
Styloclone (standalone version)	23106242	£7.99
Dual Mini LED Dice (TH LEDs, white solder mask)	08103241W	£2.99
Dual Mini LED Dice (TH LEDs, black solder mask)	08103241B	£2.99
Dual Mini LED Dice (SMD LEDs, white solder mask)	08103242W	£2.99
Dual Mini LED Dice (SMD LEDs, black solder mask)	08103242B	£2.99
Dual Hybrid Supply regulator board	18107211	£5.99
Dual Hybrid Supply front panel control board	18107212	£2.99
Dual Hybrid Supply CPU board	01106193	£3.99
Dual Hybrid Supply LCD adaptor	01106196	£2.99
MAY 2025		
Skill Tester 9000	08101241	£9.99
USB-C Serial Adaptor (black solder mask)	24106241	£2.99
DC Supply Protector (SMD adjustable version)	08106241	£3.99
DC Supply Protector (TH adjustable version)	08106242	£3.99
DC Supply Protector (TH fixed version)	08106243	£3.99
APRIL 2025		
ESR Test Tweezers PCB set	04105241+	£5.99
Fan Speed Controller Mk2	10104241	£4.99
Calibrated Microphone (SMD version)	01108231	£3.99
Calibrated Microphone (TH version)	01108232	£3.99
MARCH 2025		
Pico Gamer	08104241	£9.99
Compact Frequency Divider	04112231	£4.99
FEBRUARY 2025		
Laser Communicator transmitter	16102241	£3.99
Laser Communicator receiver	16102242	£3.99
Model Railway Points Controller transmitter	09101241	£3.99
Model Railway Points Controller receiver	09101242	£2.99
Mains Power-Up Sequencer	10108231	£11.99
JANUARY 2025		
USB to PS/2 Converter for VGA PicoMite	07111231	£3.99
ps2x2pico USB to PS/2 Converter	07111232	£3.99
Compact Microphone Preamplifier (case version)	01110231	£5.99
Compact Microphone Preamplifier (integrated version)	01110232	£5.99

PCBs for most recent *PE/EPE* constructional projects are available. From the July 2013 issue onwards, PCBs with eight-digit codes have silk screen overlays and, where applicable, are double-sided, have plated-through holes, and solder mask. They are similar to photos in the project articles. Earlier PCBs are likely to be more basic and may not include silk screen overlay, be single-sided, lack plated-through holes and solder mask.

Always check price and availability in the latest issue or online. A large number of older boards are listed for ordering on our website.

In most cases we do not supply kits or components for our projects. For older projects it is important to check the availability of all components before purchasing PCBs.

Magazine back issues are available – see https://www.electronpublishing.com/product-category/pe-back-issues-print-and-digital/

Double-sided | plated-through holes | solder mask

PROJECT	CODE	PRICE
DECEMBER 2024		
Raspberry Pi Clock Radio PCB pair	19101241/2	£12.99
VGA PicoMite (blue)	07107221	£5.99
Secure Remote Switch receiver	10109231	£4.99
Secure Remote Switch transmitter (module version)	10109232	£2.99
Secure Remote Switch transmitter (discrete version)	10109233	£2.99
NOVEMBER 2024		
Ideal Diode Bridge Rectifiers (28mm square)	18101241	£3.99
Ideal Diode Bridge Rectifiers (22mm square)	18101242	£3.99
Ideal Diode Bridge Rectifiers (4-pin SIL)	18101243	£3.99
Ideal Diode Bridge Rectifiers (tiny SOT-23)	18101244	£3.99
Ideal Diode Bridge Rectifiers (standalone SMD)	18101245	£4.99
Ideal Diode Bridge Rectifiers (standalone TH; 2oz copper)	18101246	£6.99
Multi-channel Volume Control (volume PCB)	01111221	£4.99
Multi-channel Volume Control (control PCB)	01111222	£4.99
Multi-channel Volume Control (OLED PCB, black)	01111223	£4.99
Coin Cell Emulator (black)	18101231	£3.99
OCTOBER 2024		
TQFP-32 Programming Adaptor	24108231	£4.99
TQFP-44 Programming Adaptor	24108232	£4.99
TQFP-48 Programming Adaptor	24108233	£4.99
TQFP-64 Programming Adaptor	24108234	£4.99
2m VHF FM Test Signal Generator	06107231	£4.99
SEPTEMBER 2024		
PIC Programming Adaptor	24105231	£4.99
Salad Bowl Speaker crossover	01109231	£7.99
AUGUST 2024		
Arduino ESR Meter combined L/C-ESR	04106182	£9.99
Arduino ESR Meter separate ESR	04106183	£7.99
Watering System Controller	15110231	£11.99
JULY 2024		
Dynamic NFC Tag (22 × 31mm)	06101231	£5.95
Dynamic NFC Tag (86 × 54mm)	06101232	£7.95
Reciprocal Frequency Counter	CSE230101C	£8.95
Reciprocal Frequency Counter front panel (black mask)	CSE230102	£7.95
Pi Pico-based Thermal Camera	04105231	£6.95
JUNE 2024		
Loudspeaker Test Jig	04106231	£12.95
MAY 2024		
GPS-Disciplined Oscillator	04103231	£12.95
Dual RF Amplifier	CSE220602A	£5.95
Songbird	08103231	£7.95
APRIL 2024		
Basic RF Signal Generator	CSE221001	£8.99
Basic RF Signal Generator front panel (black solder mask)	CSE220902B	£9.99
Silicon Chirp Cricket	08101231	£7.99
Test Bench 'Swiss Army Knife'	04110221	£9.99
MARCH 2024		
Digital Volume Control Potentiometer SMD version	01101231	£8.95
Digital Volume Control Potentiometer through-hole version	01101232	£9.95
Advanced SMD Test Tweezers body	04106221	£6.95
Advanced SMD Test Tweezers legs (pair)	04106212	£7.95
FEBRUARY 2024		
Active Mains Soft Starter	10110221	£10.95

PROJECT	CODE	PRICE
JANUARY 2024		
Q Meter	CSE220701	£8.95
Q Meter front panel (black solder mask)	CSE220704	£7.95
Raspberry Pi Pico W BackPack	07101221	£8.95
DECEMBER 2023		
Digital Boost Regulator	24110224	£8.95
Dual-Channel Power Supply for Breadboards	04112221	£11.95
Display Adaptor for the Breadboard PSU	04112222	£9.95
NOVEMBER 2023		
LC Meter Mk3	CSE220503C	£9.95
DC Supply Filter for vehicles	08108221	£8.95
Discrete Audio Op Amp PCB	AO1-JUL23	£9.95
Discrete Audio Op Amp PCB + essential components	AO2-JUL23	£17.95
OCTOBER 2023		
Buck/Boost Charger Adaptor	14108221	£9.95
PIC Breakout Board for SOIC parts	24110225	£5.95
PIC Breakout Board for DIP parts	24110222	£5.95
AVR64DD32 Breakout board	24110223	£5.95
Automatic Train Controller	09109221	£6.95
Chuff Sound module	09109222	£6.95
SEPTEMBER 2023		
Mini LED Driver	16106221	£7.95
New GPS-Synchronised Clock	19109221	£9.95
AUGUST 2023		
110dB RF Attenuator	CSE211003	£9.95
Universal Battery Charge Controller (2023 update)	14107192	£9.95
Wide-Range OhmMeter	04109221	£12.95
JULY 2023		
Multimeter Checker/Calibrator	04107221	£10.95
JUNE 2023		
Arduino Programmable Load	04105221	£8.95
Buck-Boost LED Driver	16103221	£8.95
MAY 2023		
Precision AM-FM DDS Signal Generator	CSE211002	£14.95
Improved SMD Test Tweezers programmed PIC	0410621PIC	£12.95
AO 2x-dual-to-quad through-hole (pack of five)	AO1-MAY23	£6.95
AO 2x-dual-to-quad SMD (pack of five)	AO2-MAY23	£6.95
AO 2x-single-to-dual through-hole (pack of five)	AO3-MAY23	£6.95
AO 2x-single-to-dual SMD (pack of five)	AO4-MAY23	£6.95
APRIL 2023		
500W Amplifier Module	see p.22, April 2023	
Clipping Indicator (per channel)	01112211	£7.95
CD Welder Power Supply (one needed)	29103221	£9.95
CD Welder Controller (one needed)	29103222	£9.95
CD Welder Energy Storage module (several needed)	29103223	£7.95
AO Universal Dual Op Amp Board	AO1-APR23	£9.95
AO Stereo RIAA precision passives kit for dual op amp	AO2-APR23	£8.95
MARCH 2023		
Pico BackPack	07101221	£9.95
Semaphore Signal (controller)	09103221	£7.95
Semaphore Signal (blade)	09103222	£5.95
FEBRUARY 2023		
Tesla Coil driver board	26102221	£9.95
Tesla Coil potentiometer board	26102222	£5.95
Cooling Fan Controller & Loudspeaker Protector	01102221	£8.95
Remote Gate Controller	11009121	£12.95

All prices include VAT and UK p&p. Add £4 per *project* for post to Europe; £5 per *project* outside Europe.

Orders and payment should be sent to:

Practical Electronics, Electron Publishing Ltd
113 Lynwood Drive, Merley, Wimborne, Dorset BH21 1UU
Tel 07918 614662   Email: pe@pemag.au
Online Shop: www.electronpublishing.com

Sorry, cheques cannot be accepted as we are based in Australia and our UK bank account does not allow cheque deposits.

NOTE: Most boards are in stock and sent within seven days of order receipt. Please allow up to 28 days delivery if we need to restock.

Double-sided | plated-through holes | solder mask

PROJECT	CODE	PRICE
JANUARY 2023		
Classic LED Metronome – 8-LED	23111211	£7.95
Classic LED Metronome – 10-LED	23111212	£8.95
Multi-Channel Speaker Protector – 6 channel	01101221	£9.95
Multi-Channel Speaker Protector – 4 channel	01101222	£7.95
Remote Control Range Extender – IR-to-UHF	15109212	£5.95
Remote Control Range Extender – UHF-to-IR	15109211	£7.95
AO Universal Single Op Amp Board	AO1-JAN23	£7.95
DECEMBER 2022		
Hummingbird Amplifier	01111211	£9.95
SMD Trainer PCB	29106211A	£8.95
NOVEMBER 2022		
USB Cable tester – main PCB	04108211	£12.95
USB Cable tester – front panel	04108212	£5.95
USB Cable tester – optional panel	SC5970	£5.95
Model Railway Carriage Lights – PCB	09109211	£6.95
AO trnsfmr PCB – standard VTX-A range	VTX-101-007	£6.95
AO trnsfmr PCB – dual-outline VTX102-3001/101-3002	VTX-Dual	£6.95
OCTOBER 2022		
SMD Test Tweezers – PCB and pair of tweezer arms	04106211-2	£11.95
SMD Test Tweezers – programmed PIC12F1572-I/SN	0410621A	£7.95
Tele-com	12110211	£12.95
SEPTEMBER 2022		
Touchscreen Digital Preamp – main board	01103191	£12.95
Touchscreen Digital Preamp – adaptor board pair	01103192	£5.95
20A DC Motor Speed Controller	11006211	£9.95
AUGUST 2022		
Multi-purpose Battery Manager – I/O Expander	11104212	£5.95
Multi-purpose Battery Manager – Switch Module	11104211	£8.95
Simple MIDI Music Keyboard (for 8 switches)	23101213	£6.95
Nano Pong	08105212	£7.95
JULY 2022		
Silicon Labs AM/FM/SW Radio	CSE210301C	£10.95
Level Crossing Controller	09108211	£6.95
JUNE 2022		
Full-wave Motor Speed Controller	1010221	£8.95
PIC Programming Helper for 8-pin PICs only	24106211	£7.95
PIC Programming Helper for 8, 14 or 20-pin PICs	24106212	£10.95
Advanced GPS Computer	05102211	£9.95
MAY 2022		
Bus board PCB for Analogue Vocoder	AO1-MAY22	£10.95
Complete set of 14 PCBs for Analogue Vocoder	AO2-MAY22	£97.95
Programmed EEPROM for Digital FX Unit	FX1-MAY22	£10.95
Programmed PIC for Digital FX Unit w/ potentiometer	FX2-MAY22	£8.95
APRIL 2022		
64-key MIDI Matrix shield	23101211	£8.95
64-key MIDI Matrix switch board	23101212	£11.95
High-current Battery Balancer	14102211	£10.95
Digital FX Unit – using potentiometer	01102211	£9.95
Digital FX Unit – using BCD switch	01102212	£9.95
Universal Audio PSU	AO1-APR22	£11.95
MARCH 2022		
Mini Isolated Serial Link	24102211	£5.95
Busy Loo Indicator	16112201	£5.95
Analogue Vocoder – Band-pass filter board	AO1-MAR22	£9.95
Analogue Vocoder – HP/LP filter board	AO2-MAR22	£9.95

PROJECT	CODE	PRICE
FEBRUARY 2022		
Arduino-based Power Supply	18106201	£9.95
Battery Monitor Logger	11106201	£10.95
Electronic Wind Chimes	23011201	£10.95
Analogue Vocoder – Driver Amplifier	AO-FEB22	£8.95
JANUARY 2022		
Vintage battery Radio Li-ion Power Supply	11111201	£9.95
MiniHeart: A Miniature Heartbeat Simulator	01109201	£8.95
DECEMBER 2021		
AM/FM/SW Digital Receiver	CSE200902A	£13.95
Balanced Input and Attenuator for USB CODEC	01106202	£11.95
NOVEMBER 2021		
Dual Battery Lifesaver	11111202	£6.95
OCTOBER 2021		
Mini Wi-Fi LCD BackPack	24106201	£8.95
SEPTEMBER 2021		
USB SuperCodec PCB	01106201	£14.95
Audio DDS Oscillator PCB	01110201	£5.95
Audio DDS Oscillator rotary encoder	01110201-ENC	6.95
Programming Adaptor Board for Audio DDS Oscillator	01110202	£5.95
High-power Ultrasonic Cleaner main PCB	04105201	£14.95
High-power Ultrasonic Cleaner front-panel PCB	04105202	

For the many earlier PCBs that we stock please see the PE website: www.electronpublishing.com

PE/EPE PCB SERVICE

Order Code	Project	Quantity	Price

Name ....

Address ....

....

Tel ....

Email ....

Card No ....

Valid From ....Expiry Date ....

Card Security No ....

You can also order PCBs by phone, email or via the shop on our website: www.electronpublishing.com

No need to cut your issue – a copy of this form is just as good!

Kits and pre-programmed microcontrollers are available from the Silicon Chip shop in Australia for many of the projects published in *Practical Electronics*.

To see a list of the kits available, visit: https://siliconchip.au/Shop/20

Pre-programmed microcontrollers are listed here: https://siliconchip.au/Shop/9
In many cases, you can also order printed circuit board(s) for the same projects: https://siliconchip.au/Shop/8

VAT is charged to UK customers at check-out from the Silicon Chip shop and currently, orders under £85 should not attract any extra duties.
Please allow 1-2 weeks for delivery (longer in some cases).

CLASSIFIED ADVERTISING

Practical Electronics

If you want your advertisements to be seen by the largest readership at the most economical price then our classified page offers excellent value.

The rate for semi-display space is £10 per centimetre high, with a minimum height of 2·5cm. All semi-display advertisements have a width of 5.5cm. The prepaid rate for classified adverts is 40p per word (minimum 12 words).

For rates, further information on display and classified advertising or to book an ad, please email us at pe@pemag.au

Practical Electronics reaches more UK readers than any other UK monthly hobby electronics magazine. Our sales figures prove it.

We have been the leading monthly magazine in this market for more than thirty years.

W5SWL Electronics

Technical Products and Small Component Parts, RF and Audio Connectors, Hobby Electronics, Ham Radio to Microwave Products, General Hobby Supplies, Technical Books, Tools, and More. **Located in the United States.**

Visit **www.W5SWL.com**

BOWOOD ELECTRONICS LTD

Suppliers of Electronic Components

www.bowood-electronics.co.uk

Unit 10, Boythorpe Business Park, Dock Walk, Chesterfield, Derbyshire S40 2QR. Sales: 01246 200 222

Send large letter stamp for Catalogue

COAST ELECTRONICS

BREAKOUTS-COMPONENTS-CONTRACT DESIGN-3D PRINTER PARTS-MUSICAL-MICROCONTROLLERS

WWW.COASTELECTRONICS.CO.UK

MISCELLANEOUS

VALVES AND ALLIED COMPONENTS? For free stock list and/or advice, please contact me: geoffdavies337@gmail.com Telephone: 01788 574774

PIC DEVELOPMENT KITS, DTMF kits and modules, CTCSS Encoder and Decoder/Display kits. Visit www.cstech.co.uk

Electrical Industries Charity (EIC)

We help people working in the electrical, electronics and energy community as well as their family members and retirees.

We use workplace programmes that give the industry access to financial grants and a comprehensive range of free and confidential services.

www.electricalcharity.org

ADVERTISING INDEX

AOSHOP	19
ENIGMA COMPONENT SHOP	9
MICROCHIP TECHNOLOGY	Outer back cover
PCBWAY	Inner front cover
PDF BACK ISSUE DOWNLOADS	35 & 74
PO LABS	77
TEACH-IN BUNDLES & DVDS	18, 55 & 79

Advertising mailing address

Electron Publishing Australia
PO Box 194
Matraville NSW 2036
Australia

Email pe@pemag.au

Web www.electronpublishing.com

For editorial contact details, see page 2.

Next Month – in the August issue

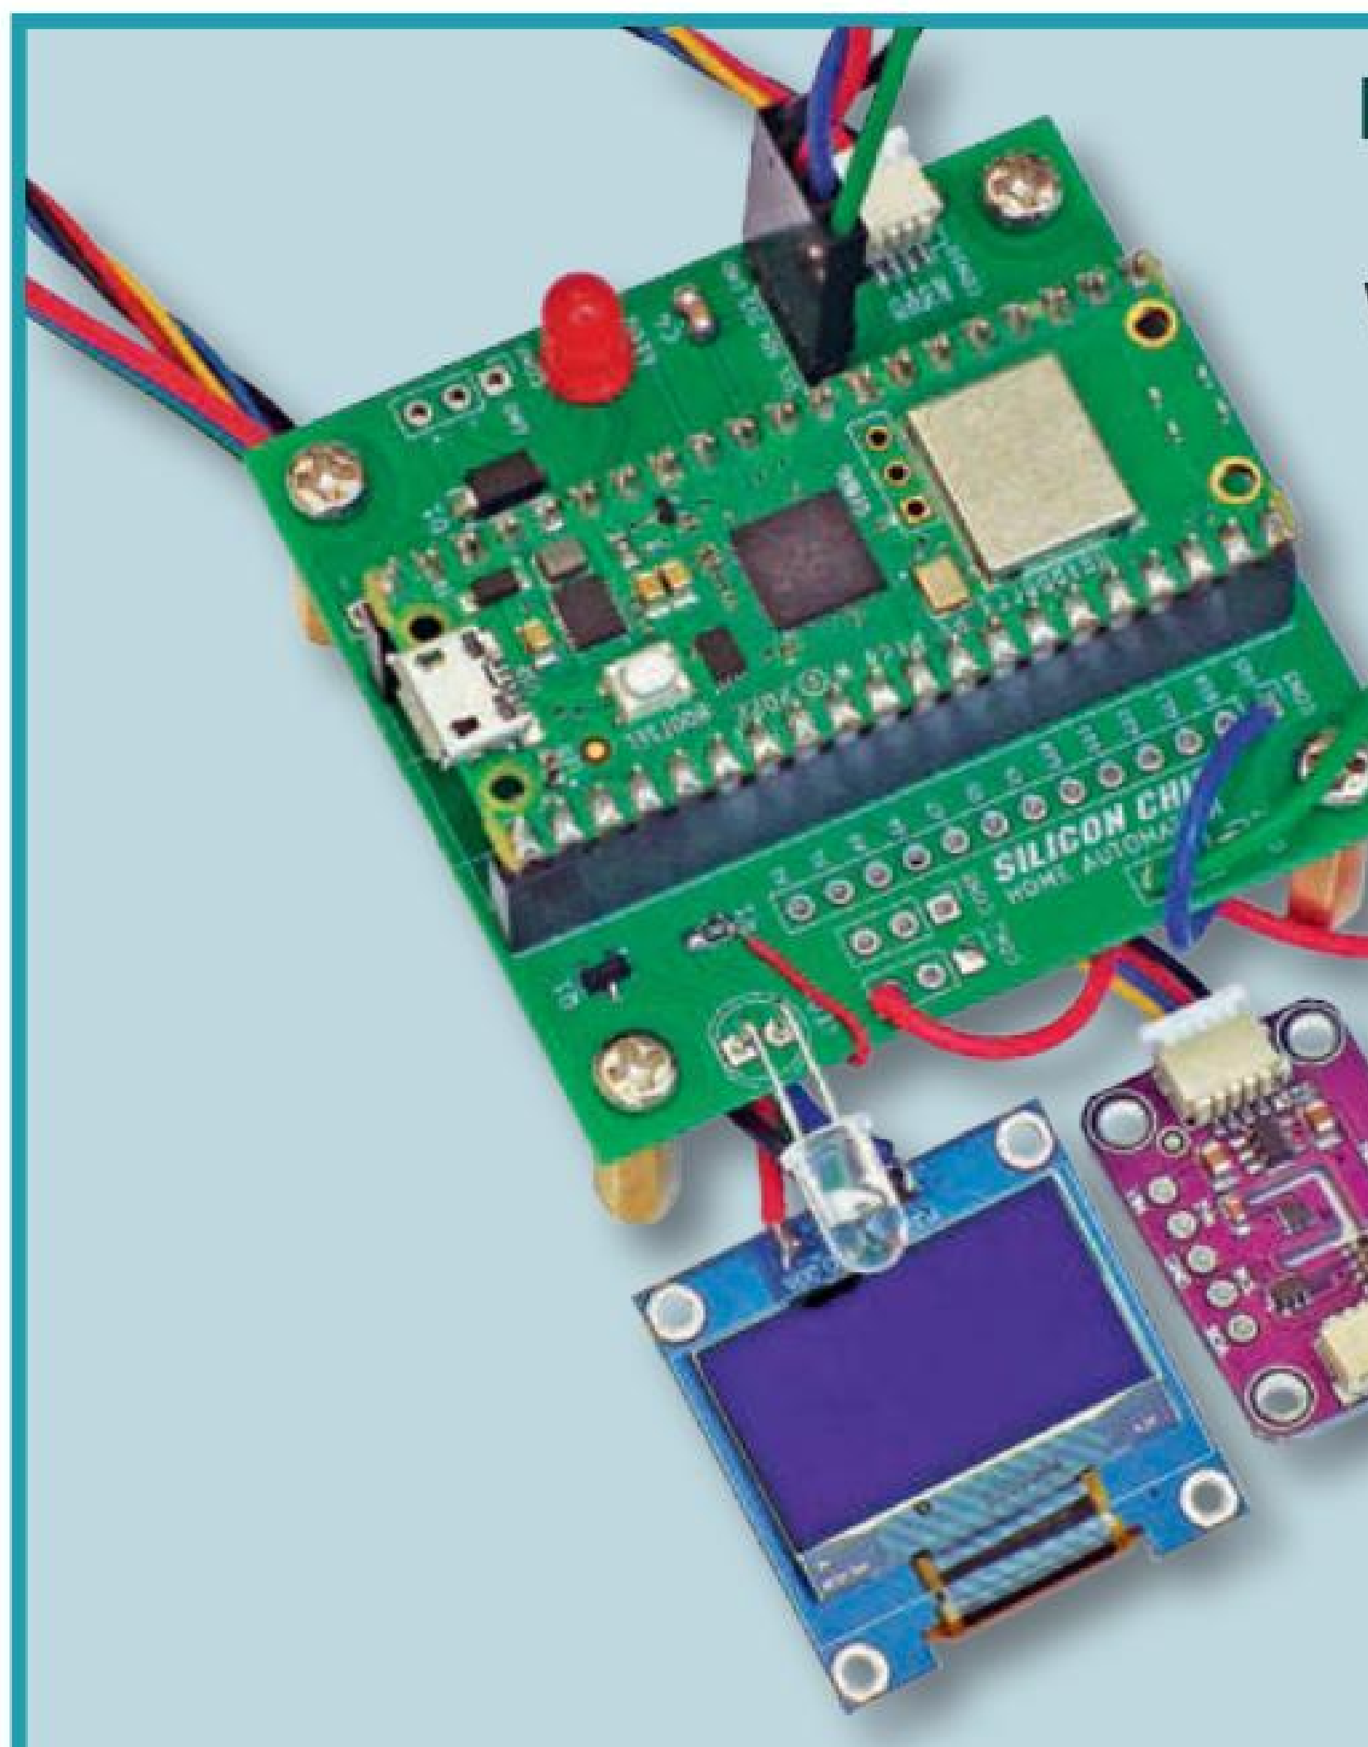

Home Automation Satellite

We explain how to set up a HomeAssistant hub, then connect sensors and displays

4.86V -> 0.472A
+2.293W START
0.000J RESET
CHG
USB PWR
SILICON CHIP
USB-C Power Monitor

USB-C Power Monitor

Measures voltage, current, power and energy with bidirectional and USB PD support

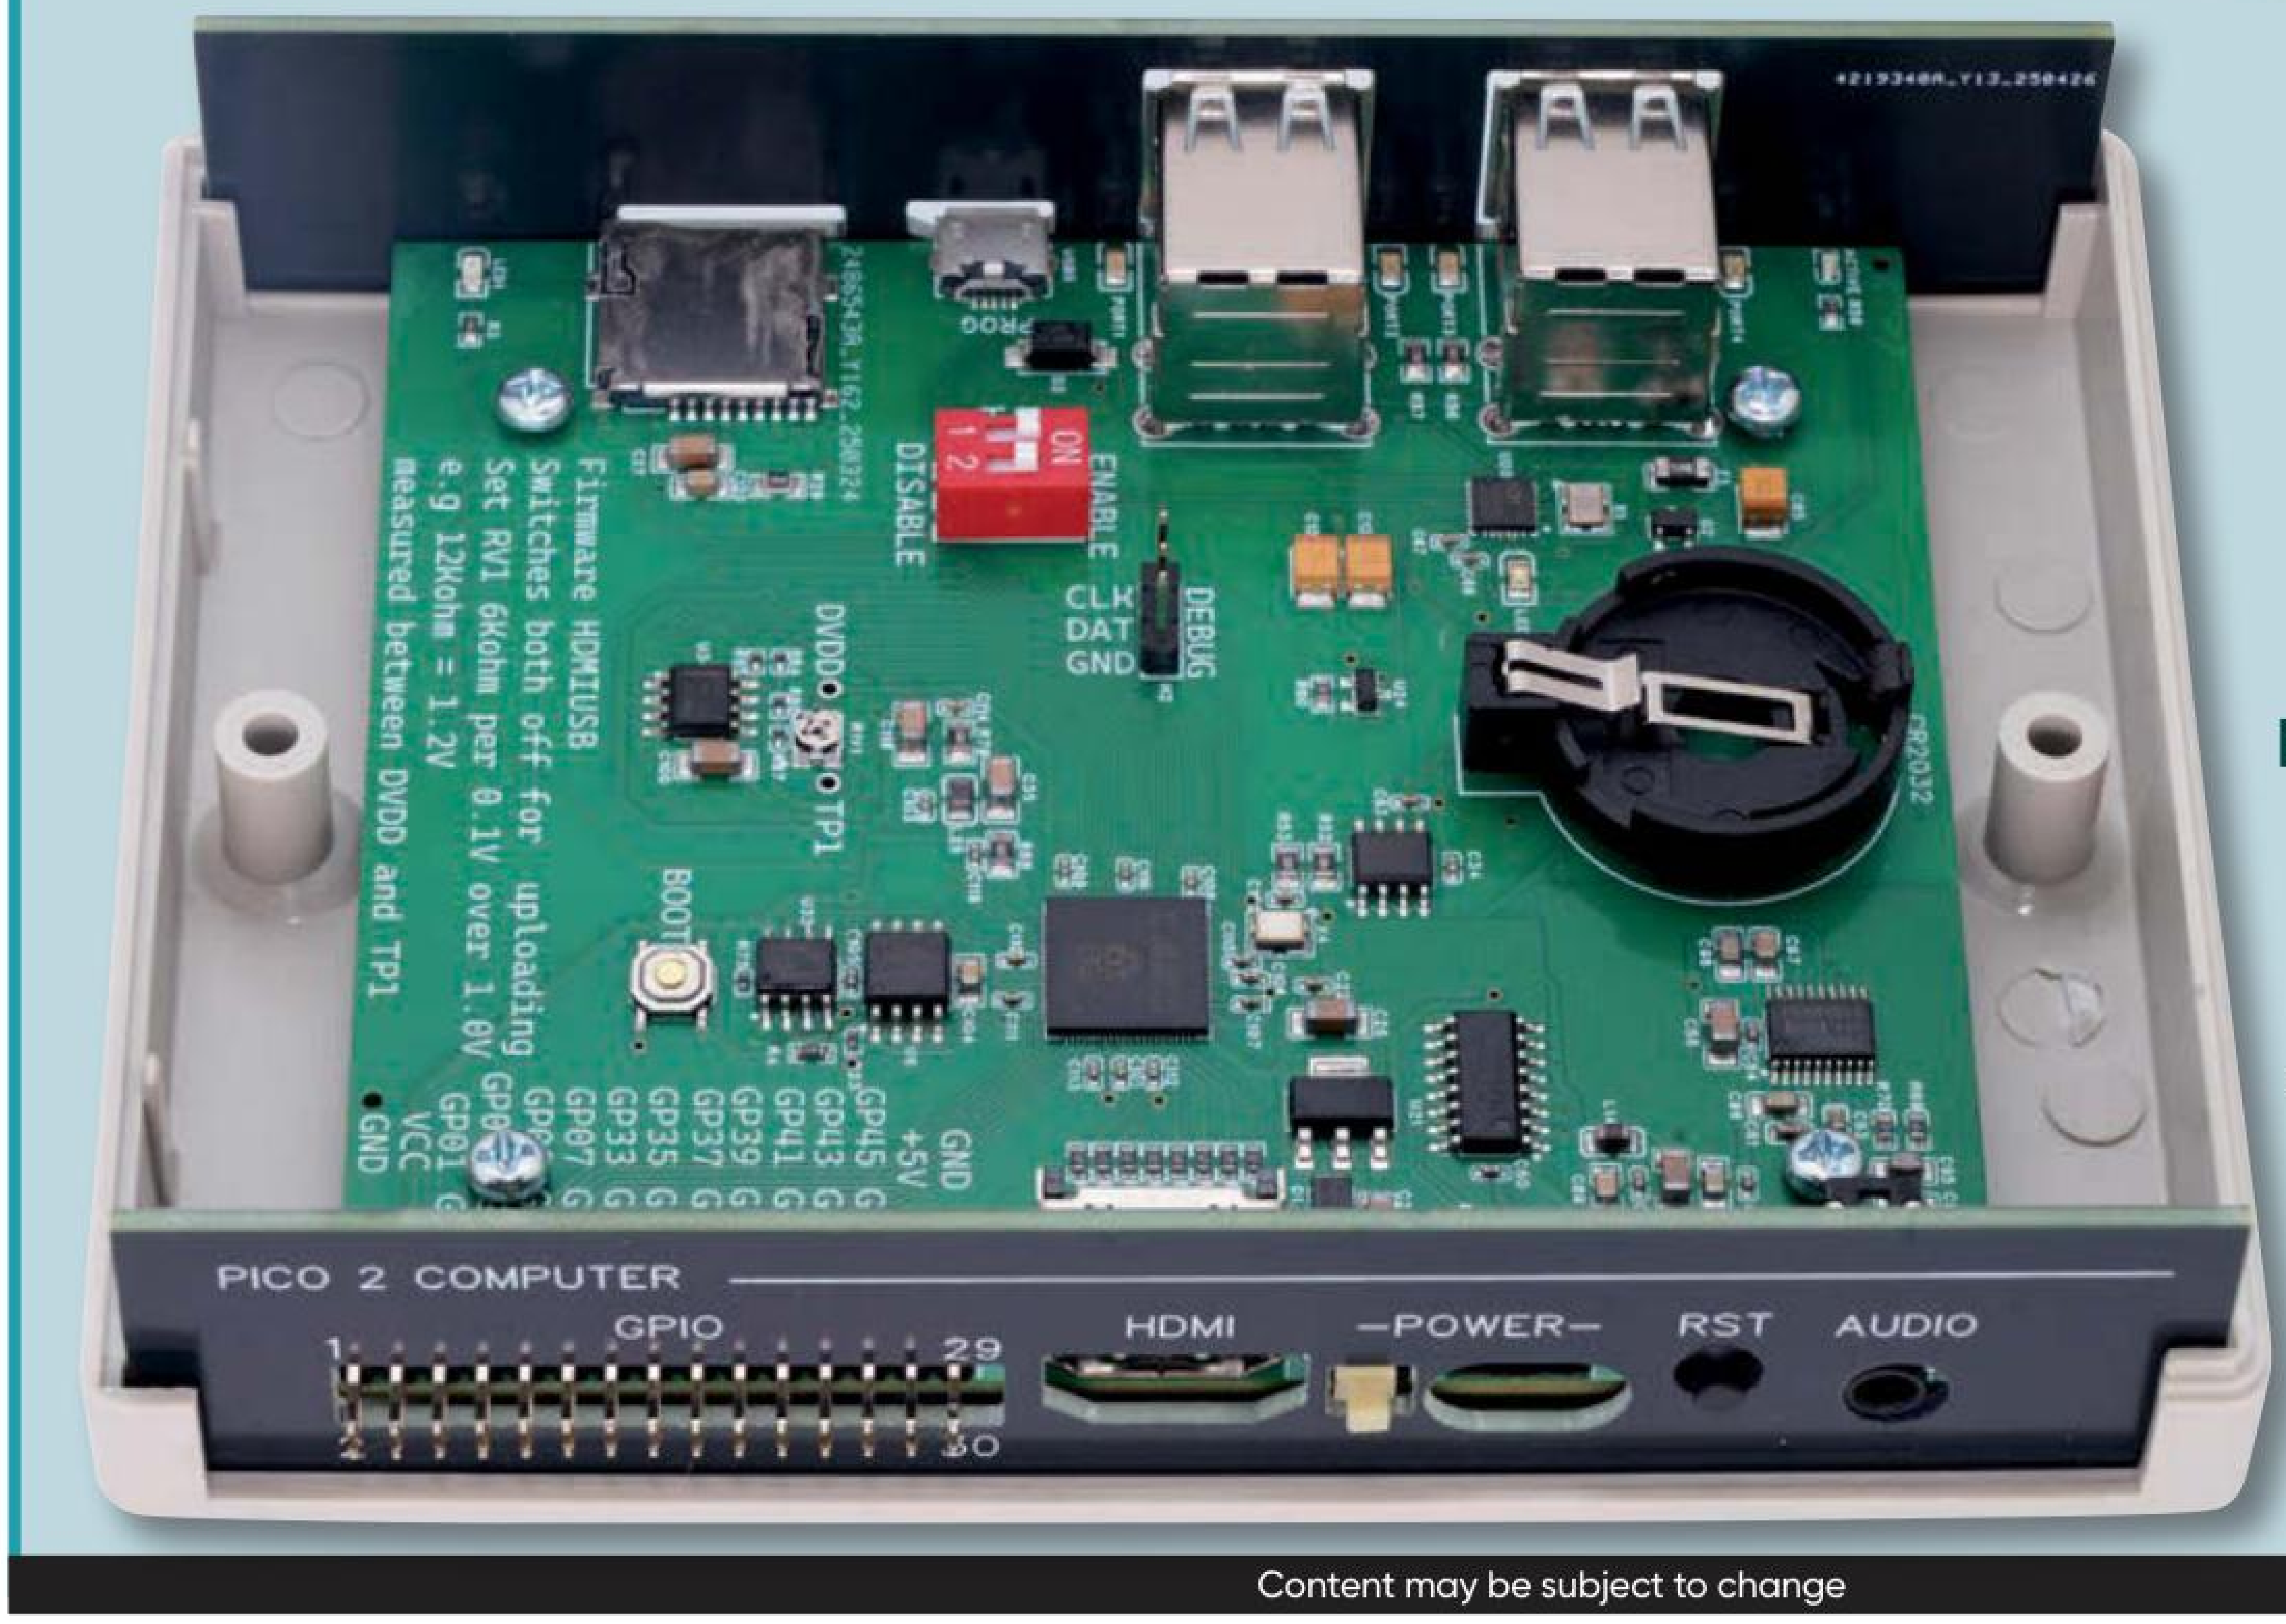

RP2350B Computer

HDMI video, four USB ports for keyboards, mice and game controllers, stereo audio and expandable storage/memory

On sale 7th of July, 2026

Content may be subject to change

Practical Electronics

Subscriptions hotline

To subscribe or renew online, paying by credit/debit card or PayPal, please visit our shop at: www.electronpublishing.com

For those wishing to pay by direct debit, cheque or over the telephone, please call 'Select Publisher Services' on 01202 087631, or send a cheque (payable to: 'Select Publisher Services Ltd') with your details to:

Practical Electronics Subscriptions, PO Box 6337, Bournemouth BH1 9EH, United Kingdom

Remember, we print the date of the last issue of your current subscription in a box on the address sheet that comes with your copy.

Digital subscribers, please call 07918 614662 or visit: www.electronpublishing.com

Published on approximately the first Thursday of each month by Silicon Chip Publications, PO Box 194 Matraville NSW 2036 Australia. Printed in England by The Manson Group, Reynolds House, 8 Porters Wood, Valley Road Industrial Estate, St Albans AL3 6PZ. Distributed by Seymour, 86 Newman St, London W1T 3EX. Subscriptions UK: £33.99 (6mo); £59.99 (12mo.); £114.99 (2yrs). EUROPE: airmail service, £38.99 (6mo); £69.99 (12mo); £129.99 (2yrs). REST OF THE WORLD: airmail service, £46.99 (6mo); £84.99 (12mo); £164.99 (2yrs). Make payments to 'Practical Electronics', Practical Electronics Subscriptions, PO Box 6337, Bournemouth BH1 9EH United Kingdom. Email: pesubs@selectps.com. PRACTICAL ELECTRONICS is sold subject to the conditions that it shall not, without the written consent of the Publishers first having been given, be lent, resold, hired out or otherwise disposed of by way of Trade at more than the recommended selling price on the cover, and that it shall not be lent, resold, hired out or otherwise disposed of in a mutilated condition or in any unauthorised cover by way of Trade or affixed to or as part of any publication or advertising, literary or pictorial matter.

Did you know our online shop now sells the current issue of *PE* for £5.99 inc. p&p?

Practical Electronics
The UK's premier electronics, computing and maker magazine
Max's Cool Beans
Playing games by tilting our console
Teach-In 12
Part 9: RF sources and oscillators
Audio Out
How symmetricisers boost radio volume
Circuit Surgery
Sending & receiving balanced audio
RP2350B Dev Board: A more powerful Pico 2
Audio Out: a flexible all-pass filter system
The SmartProbe: make measurements more easily
PLUS!
Techno Talk – the history of PCB design and track routing
The Fox Report – balcony solar may be coming to the UK
Mains Earthing – why it's needed & how it varies worldwide
Jul 2026 £5.99
www.electronpublishing.com
@practicalelec
practicalelectronics
Audio Out
An 8-way All-pass Filter PCB
Circuit Surgery
Differential and balanced signals
A Simple USB Power Adaptor
DIY 433MHz Receiver
Jun 2026 £5.99

You read that right! We now sell the current issue of your favourite electronics magazine for exactly the same price as in the High Street, but we deliver it straight to your door – and for UK addresses we pay the postage. No need to journey into town to queue outside the newsagent. Just go to our website, set up an account in 30 seconds, order your magazine and we'll do the rest.

www.electronpublishing.com

PIC18-Q35: Processing Power Meets Programmable Logic

Zero CPU Overhead, Maximum Flexibility

The PIC18-Q35 microcontroller family empowers you to design smarter, more efficient embedded systems by combining high-speed processing with advanced, on-chip programmable logic. The integrated Configurable Logic Block (CLB) with 128 Basic Logic Elements lets you create custom peripherals and timing functions, no external logic ICs required.

Key Features

- **Custom Logic, Zero CPU Load:** CLB offloads real-time control for deterministic performance
- **Instant, Reliable Operation:** Logic auto-loads from Flash at every boot
- **Simplified, Secure Design:** Eliminate external logic, reduce board space and minimize the risk of software-based attacks
- **Efficient data transfers:** with Direct Memory Access (DMA)
- **Comprehensive suite of advanced analog and digital peripherals:** to support a variety of applications without requiring extra components
- **User-friendly GUI:** that simplifies peripheral and CLB configuration

Solve Real-World Challenges

- Remove CPU bottlenecks and maintain real-time performance
- Simplify code, no more complex boot-time configuration routines
- Configure the exact peripheral you need, right on-chip

Free your CPU from real-time control and achieve true deterministic performance, even in the most demanding industrial, automotive and IoT applications.

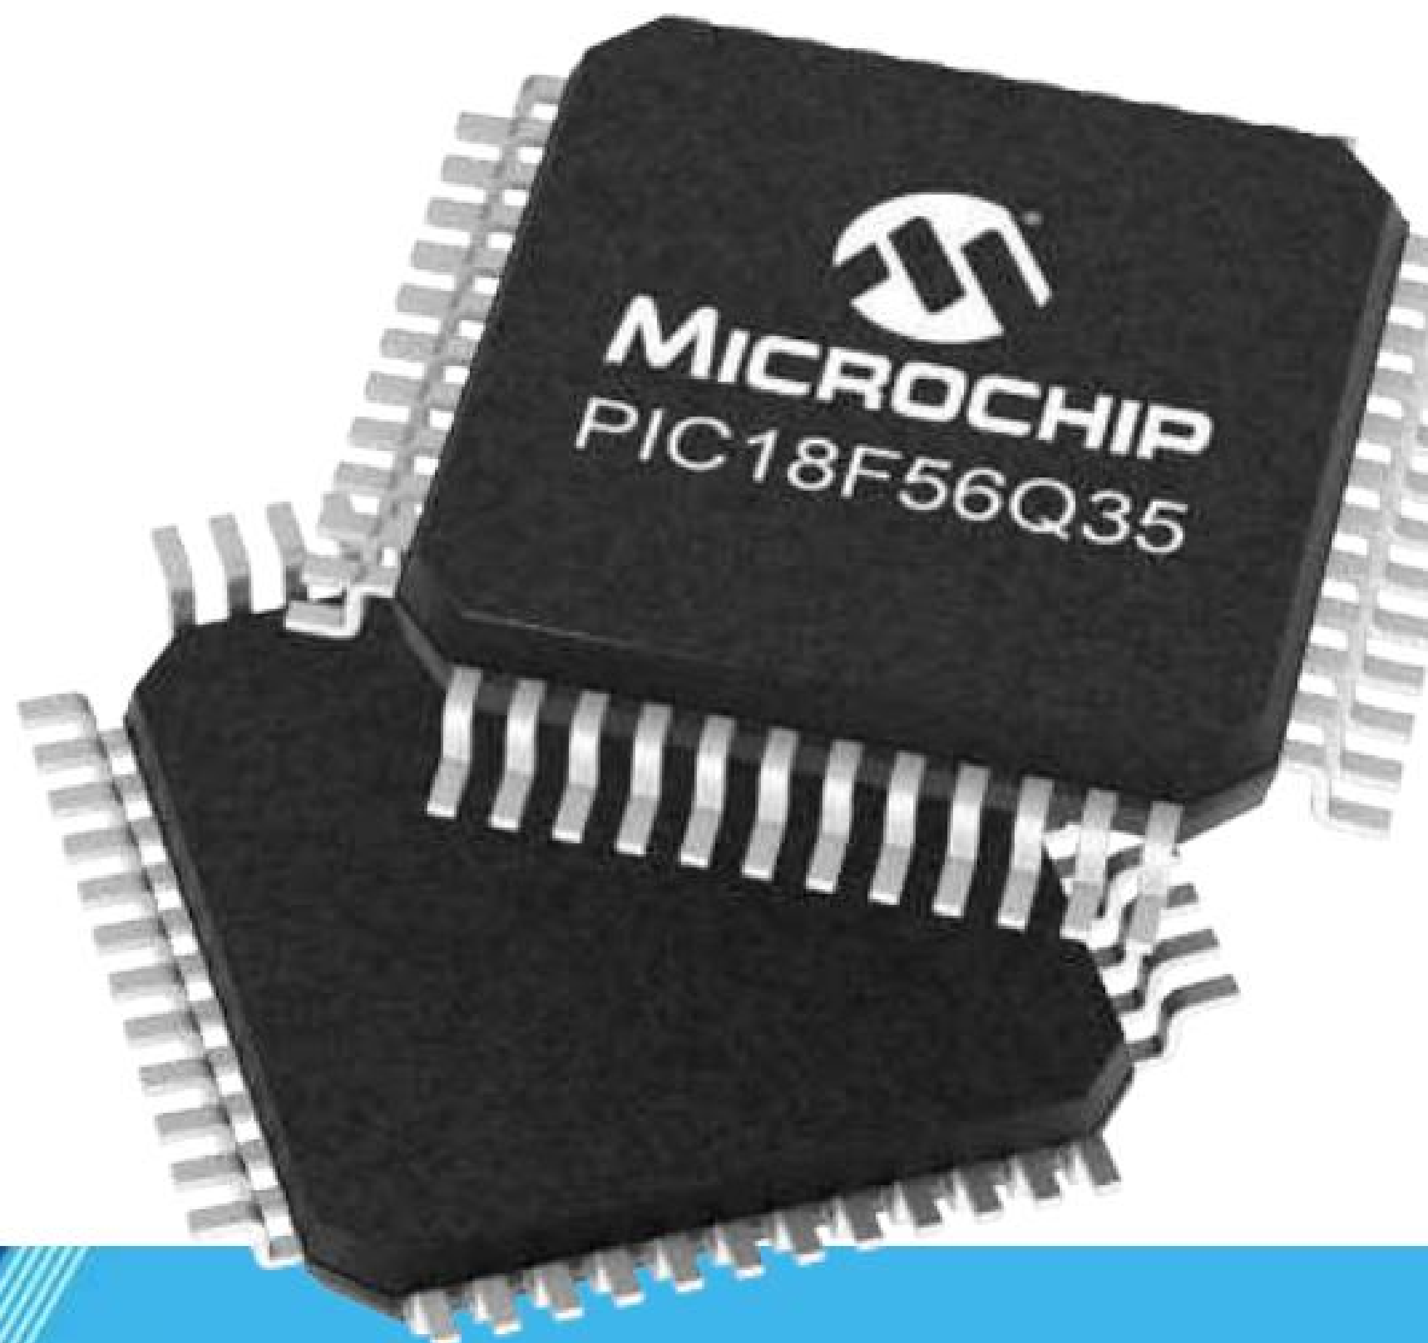

MICROCHIP

microchip.com/pic18-q35

The Microchip name and logo and the Microchip logo are registered trademarks of Microchip Technology Incorporated in the U.S.A. and other countries. All other trademarks are the property of their registered owners.
© 2026 Microchip Technology Inc. All rights reserved.
MEC2649A-UK-05-26